# Reachability in Vehicle Lateral Stability

## Application of Hamilton-Jacobi Safety Filter as Electronic Stability Control

Bartosz Jemioł

TU Delft

# Reachability in Vehicle Lateral Stability

## Application of Hamilton-Jacobi Safety Filter as Electronic Stability Control

Thesis report

by

Bartosz Jemioł

to obtain the degree of Master of Science
at the Delft University of Technology
to be defended publicly on November 12, 2025 at 13:00

*Thesis committee*:

| | |
|---|---|
| Chair: | Dr.ir. D.M. Pool |
| Supervisor: | Dr.ir. C.C. de Visser |
| External examiner: | Dr. B. Shyrokau |
| Place: | Faculty of Aerospace Engineering, Delft |
| Project Duration: | January, 2025 - November, 2025 |
| Student number: | 5275830 |

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

Faculty of Aerospace Engineering · Delft University of Technology

# Acknowledgments

This thesis would not have been possible without the support and encouragement of many people who made this journey both challenging and rewarding.

First and foremost, I would like to thank my supervisors, Prof. Coen de Visser and Prof. Barys Shyrokau, for creating an environment where creativity was not only permitted but actively encouraged. Their guidance helped me balance ambitious ideas with practical implementation, and their constructive feedback was instrumental in shaping this work.

To the ”Aquarium”—my friend group who knows who they are—thank you for making this entire master’s program, and especially the thesis period, an unforgettable experience. Your companionship, humor, and support turned what could have been a stressful time into something I will always look back on with fondness.

Finally, I am deeply grateful to my parents for their unwavering belief in me and for providing me with opportunities that have shaped who I am today. They have given me an education and a student life that I don’t think could have been any better.

*Bartosz Jemioł*
Delft, November 2025

# Contents

**1 Literature Review** **1**
1.1 Hamilton-Jacobi Reachability and Its Applications . . . 2
1.2 Safety Filter Approaches for Lateral Stability in Vehicles . . . 3

**2 Project Plan** **6**
2.1 Research Gap and Proposed Approach . . . 6
2.2 Research Question(s) . . . 7
2.3 Methods, Tools, and Expected Results . . . 7
2.4 Planning . . . 9

**I Scientific Article** **13**

**3 Hamilton-Jacobi Reachability for Vehicle Lateral Stability: A Control Barrier-Value Function Approach** **14**
3.1 Introduction . . . 14
3.2 Background . . . 15
3.3 Methodology . . . 18
3.4 Results . . . 22
3.5 Conclusion & Future Work . . . 33

**II Additional Results** **36**

**4 Additional Testing** **37**
4.1 Industry Controller Performance Analysis . . . 37

**5 Verification** **40**
5.1 HJ Solver . . . 40
5.2 Quadratic Programming Controller . . . 41

**6 Conclusion** **46**
6.1 Closing Remarks . . . 46
6.2 Research Questions . . . 46
6.3 Scientific Contributions . . . 47

**7 Recommendations** **49**

**References** **52**

# 1

# Literature Review

Safety is a critical concern in vehicle control systems, particularly for autonomous and semi-autonomous vehicles. Recent research has focused on developing frameworks that guarantee vehicle stability while allowing for maximum operational freedom. This literature review explores the current state of research in safety envelope computation for vehicle systems, with a particular focus on lateral stability control and the concept of controlled invariant sets. The review also highlights a significant research gap in applying Hamilton-Jacobi (HJ) reachability analysis to vehicle lateral stability problems, which presents a promising opportunity for advancing the state of the art in vehicle safety systems.

Recent years have seen significant progress in safety-critical control approaches for autonomous systems. The field has come to call these controllers ”safe filters.” Safety filters monitor a system's operation at runtime and intervene when necessary to prevent failures [1]. These safety filters decouple safety from performance objectives, allowing for independent design of safety-agnostic task policies and task-agnostic safety filters.

The safety filter concept has been applied across various domains in the field. Hamilton-Jacobi reachability analysis computes the maximal controlled invariant set through dynamic programming techniques, providing a mathematical foundation for safety guarantees. Control Barrier Functions (CBFs) offer an alternative approach by ensuring safety through a Lyapunov-like derivative condition that creates a barrier between safe and unsafe regions of operation [2, 3]. Additionally, Model Predictive Safety Filters verify safety at runtime through forward simulation of the system dynamics, enabling preemptive intervention before safety violations occur [4].

In the domain of vehicle dynamics, lateral stability is particularly critical as it relates to the vehicle's ability to maintain control during cornering and lane-change maneuvers. Looking at a comprehensive survey of lateral stability criteria and control applications for autonomous vehicles, two primary methods emerge [5]. Traditional linear vehicle stability analysis relies on steering characteristics, categorizing vehicles as either understeer or oversteer. More sophisticated nonlinear analysis techniques employ phase plane methods to visualize stability regions across different state variables. These include $\beta - r$, the $\beta - \dot{\beta}$, and the $\alpha_f - \alpha_r$ phase planes. These phase plane analyses enable visualization of stability regions and determination of conditions under which instability occurs.

The visualizations do provide insight into the regions of stability, but they do not provide mathematical guarantees of stability. There are other methods used, like Lyapunov function-based analysis or Koopman Operator's spectral analysis, to give mathematically calculated safety regions, but these approaches come with their own limitations [6, 7]. Lyapunov-based approaches necessarily look at infinite time horizons, since that is how Lyapunov stability is defined. Koopman spectral analysis can provide safety regions for limited time domains; however, a limitation that it shares with Lyapunov analysis is the complex nature of the derivation.

Hamilton-Jacobi reachability analysis does provide mathematical guarantees for finite time, and it does not need a complex derivation to be used. HJ reachability is meant to deal with non-linear and complex systems, therefore, it should work well with the vehicle lateral stability problem. Additionally, with a single exception, it has not been used in lateral vehicle control [8]. The combination of a suited nature and the unexplored usage provides a strong motivation for the thesis.

## 1.1. Hamilton-Jacobi Reachability and Its Applications

Hamilton-Jacobi reachability has been studied extensively. Comprehensive overviews of HJ reachability, discussing the mathematical formulation based on dynamic programming, computation of the maximal controlled invariant set through the HJ partial differential equation, and recent advances in computational techniques have been presented [2]. Extensions that include backward reachable sets for safety verification, synthesis of optimal safety policies and applications, including aircraft collision avoidance, have been demonstrated as well [9]. Recent work also included using HJ reachability as a means for safe learning, where the safe set is gradually expanded, as the model uncertainty decreases through learning [10].

Particularly noteworthy for vehicle applications is the work on database-driven approaches to safe envelope prediction [11]. An innovative database-driven system called DEFEND (Database-Driven Safe Flight Envelope Prediction System) has been developed, which addresses the computational challenges of online reachability analysis. Their approach precomputes safety envelopes for different operating conditions and stores them in an onboard database, allowing for real-time retrieval during operation. When the required safety envelope cannot be directly retrieved from the sparse database, a novel fast-marching-based interpolation scheme has been proposed to approximate intermediate envelopes with high accuracy.

This method offers significant computational advantages, demonstrating speed-up ratios exceeding 40,000 for two-dimensional problems and over $10^6$ for higher-dimensional systems compared to direct, online computation. These approaches are particularly promising for vehicle applications where computational resources are limited. The interpolation technique works by establishing optimal paths between points on the boundaries of known safe sets and then performing linear interpolation along these paths to determine the intermediate safe set boundary.

However, while this work provides a practical framework for interpolating between safe sets, a significant open question remains regarding formal safety guarantees for these interpolated sets. The empirical validation shows good approximation quality, but does not provide mathematical proof that the interpolated sets preserve the controlled invariance properties of the original sets. This theoretical gap represents an important research opportunity: developing a framework that not only enables efficient interpolation between safe sets computed via HJ reachability methods but also provides formal guarantees that the resulting interpolated sets maintain the necessary safety properties for vehicle control applications.

Recent years have seen significant progress in developing formal methods for safe trajectory planning and control of autonomous systems. Two example frameworks that used HJ reachability as part of their control strategy are FaSTrack and PARC [12, 13]. FaSTrack was developed first and addresses a fundamental challenge in safe autonomous navigation: achieving both real-time planning capability and guaranteed safety. The framework achieves this by decoupling the planning and safety aspects of autonomous navigation, allowing any trajectory planner to use a simplified ”planning model” of the system for real-time computation, while representing the actual system with a more realistic, higher-dimensional ”tracking model.”

At the core of FaSTrack is the precomputation of a Tracking Error Bound (TEB) that accounts for the mismatch between the two models and for external disturbances. This is accomplished by modeling the navigation task as a pursuit-evasion game between the sophisticated tracking model (pursuer) and the simplified planning model (evader). The precomputation is done offline using Hamilton-Jacobi reachability analysis, yielding both the TEB and a corresponding optimal tracking controller. During online operation, obstacles are augmented by the TEB, ensuring that any potentially unsafe paths cannot be computed. The system then uses any chosen planning algorithm with the simplified model to determine the next planning state. The autonomous system applies the optimal tracking controller when near the boundary of the TEB, while any tracking controller may be used when far from violation.

Building on the foundation laid by FaSTrack, PARC (Piecewise Affine Reach-avoid Computation) was developed. It addresses the FaSTrack framework’s conservativeness while maintaining safety guarantees. PARC specifically focuses on the challenging ”near danger” scenario where a system must navigate through tightly spaced obstacles to reach a goal that lies close to these obstacles. PARC builds on two key insights to effectively solve this problem. First, it adopts a planner-tracker framework that treats the planner and tracking controller as cooperative rather than adversarial, reducing conservativeness compared to approaches like FaSTrack. Second, by representing the motion planner as a piecewise

affine (PWA) system, PARC tightly approximates the reachable set with minimal numerical approximation error. Rather than using a worst-case error bound like FaSTrack, PARC incorporates time-varying tracking errors that consider the temporal correlation of these errors, significantly reducing conservativeness while maintaining safety guarantees. The authors demonstrate this through experiments showing that PARC consistently outperforms state-of-the-art approaches, including FaSTrack, in near-danger scenarios. Perhaps most impressively, PARC enables computation of provably safe drift parking trajectories for autonomous vehicles—a task where the vehicle temporarily loses controllability. This application extends the state of the art in reachability-based safe motion planning to extreme vehicle dynamics that were previously challenging to handle with formal methods.

The field of safety envelope computation using Hamilton-Jacobi reachability methods has advanced significantly, addressing key challenges of computational complexity, real-time implementation, and formal safety guarantees. From database-driven approaches that enable efficient retrieval and interpolation of pre-computed safety envelopes to modular frameworks like FaSTrack and PARC that balance planning efficiency with safety guarantees, these methods offer promising pathways for practical implementation in autonomous vehicles. The evolution from static, conservative safety bounds to time-varying, context-aware envelopes demonstrates the field's progression toward handling increasingly complex scenarios, including extreme maneuvers like drift parking. Despite these advances, several open research questions remain, particularly regarding formal guarantees for interpolated safety envelopes and extension to higher-dimensional system models.

## 1.2. Safety Filter Approaches for Lateral Stability in Vehicles

A comprehensive review of state-of-the-art vehicle stability criteria and control applications specifically for autonomous vehicles has been written [5]. This survey systematically elaborates on the establishment, classifications, and influential factors of vehicle dynamical stability criteria, with particular emphasis on lateral stability, which is identified as "the most common and dangerous" type of vehicle instability. The survey highlights that approximately 10% of fatal crashes in the United States in 2015 were attributed to vehicle instability, underscoring the critical importance of lateral stability for driving safety.

The survey examines both linear and nonlinear stability analysis approaches. While earlier research relied on steering characteristics (understeer, neutral, and oversteer) as stability criteria, it is demonstrated that such linear criteria are insufficient since actual vehicle instability typically occurs within the nonlinear operating range of tire forces. A thorough case is made for the phase plane method as a powerful technique for analyzing nonlinearities in vehicle stability problems, exploring various implementations, including the portrait phase method and Lyapunov Function method (LFM) for estimating stability regions.

The authors identify four main phase plane representations that have been widely employed in lateral stability research: yaw rate-lateral velocity ($v_y - r$), sideslip angle-yaw rate ($\beta - r$), sideslip angle-sideslip derivative ($\beta - \dot{\beta}$), and tire slip angles ($\alpha_f - \alpha_r$) planes. Each of these representations emphasizes different aspects of vehicle dynamics and offers distinct advantages for stability analysis. The $v_y - r$ phase plane, first applied to vehicle stability in 1977, represents the direct lateral motion and rotation rate of the vehicle [14]. While this representation is straightforward, it has lower relevance to vehicle stability compared to other phase planes and is more sensitive to system parameter changes. The $\beta - r$ phase plane, which has become a primary representation in vehicle stability research, offers better relevance to stability by incorporating the sideslip angle, which, according to the survey, is a more critical indicator of vehicle stability using the "$\beta$-method" [15]. The $\beta - \dot{\beta}$ phase plane, introduced in 1994 [16], has the notable advantage that its saddle point is always located on the horizontal axis, making stability region boundaries easier to construct. This representation also demonstrates less sensitivity to parameter changes compared to other phase plane types, offering more robust stability assessment across different driving conditions. Various methods for defining stability regions in this plane have evolved, including the "two-line method" and the "diamond stability region" approach. The $\alpha_f - \alpha_r$ phase plane, which represents the slip angles of the front and rear tires, is most effective at analyzing the saturated tire force characteristics, as vehicle instability typically occurs in the nonlinear range of tire forces. This approach was first used in 1986 with the "handling-diagram method" and allows detailed analysis of the bifurcation characteristics in understeer and oversteer vehicles [17].

The survey provides a comparative analysis of these phase plane types based on three key criteria: measurability (how easily the state parameters can be measured or estimated), relevance (the degree to

which the state parameters relate to vehicle stability), and sensitivity (how the stability region changes with system parameters and disturbances). Based on this comparison, the conclusion is that the $\beta - r$ and $\beta - \dot{\beta}$ phase planes are more suitable for vehicle stability research. While the survey presents a wide variety of methods for creating safe boundaries, most of them are based on visual interpretations of graphs. The only exception presented are Lyapunov function-based methods.

Lyapunov-based approaches to lateral stability analysis of vehicle systems using nonlinear polynomial sum-of-squares (SOS) programming have shown promising results [6]. The approach focuses on determining the largest state-space region in which lateral stability can be guaranteed using peak-bounded control inputs. A computation technique is developed to calculate the maximum controlled invariant set of vehicle systems. This approach involves a three-step iterative method: first determining the region of attraction of the uncontrolled system, then verifying a local control Lyapunov function, and finally checking the acceptability and possible enlargement of the controlled invariant set. The method is computationally efficient and effectively handles the nonlinearities in vehicle dynamics.

A distinguishing feature of this approach is its ability to handle a non-linear polynomial tire model, which captures the nonlinear characteristics of lateral tire forces within a wide operation range (approximately $\pm 12^{\circ}$ of side slip). This polynomial approximation allows for more accurate stability analysis compared to conventional linear tire models, which are only valid in a narrow tire side-slip range. The vehicle model is transformed into a polynomial state-space representation where the states are the tire slip angles of the front and rear axles ($\alpha_f$ and $\alpha_r$). The choice of the $\alpha_f - \alpha_r$ phase plane, for the invariant set calculation, made sense from the perspective of implementing the nonlinear tire model. However, as explained earlier, estimating tire slip angle is the most challenging vehicle state to predict. So, for a real-life application, these calculations would need to be done using a different phase plane.

The research provides computational results showing how the maximum controlled invariant sets vary with different vehicle speeds, road adhesion coefficients, and control actuation. Their findings demonstrate that both velocity and road conditions influence the stability regions that can be guaranteed. For example, at higher velocities, smaller regions can be stabilized with the same peak-bounded control inputs, indicating that safety margins decrease as speed increases.

However, a noteworthy limitation in the approach is the unusual coupling between the maximum steering angle and steering rate. This simplification was introduced to reduce the model's complexity by avoiding a second-order steering model, but it leads to a counterintuitive result: a more capable actuator (one with faster steering capabilities and a higher maximal steering angle) produces a more constrained safe envelope. This contradicts practical intuition that greater actuation capability should expand, not contract, the region where stability can be guaranteed. By forcing faster actuators to always use their maximum rate capabilities rather than allowing them the flexibility to use gentler inputs when appropriate, the model's practical applicability is limited.

Another limitation lies in the use of Lyapunov stability theory, which is concerned with asymptotic convergence over an infinite time horizon. While this mathematically guarantees that the system will eventually return to a stable state, it provides no bounds on the recovery time. In practical vehicle dynamics scenarios, especially safety-critical situations, the time required to return to a safe operational window is crucial. A system that theoretically converges to stability but requires an impractically long time to do so may still result in a crash or loss of control in real-world conditions. The paper does not address this temporal dimension of stability, which limits its practical relevance for emergency scenarios where rapid recovery is essential.

The paper includes a validation section with simulation results comparing controlled and uncontrolled vehicles during an obstacle avoidance maneuver. The control strategy is a bang-bang approach based on the current lateral slip values, but there is limited discussion on how the system behaves during transitions or how robustly it handles operating points near the boundaries of these sets. Despite these limitations, the paper concludes with a demonstration of how the computed maximum controlled invariant sets can be implemented in a vehicle control system to improve stability during emergency maneuvers. Their simulation results confirm that intervention based on the invariant set analysis can effectively prevent vehicle instability and road departure during critical situations, though the temporal performance considerations remain unaddressed.

Another approach, based on Koopman operator analysis, aims to eliminate some of the limitations of Lyapunov-based methods [7]. It identifies lateral stability regions in vehicle dynamics using Koopman

spectral theory. While Lyapunov approaches focus on controlled invariant sets using polynomial SOS programming, the presented method leverages the Koopman operator's spectral properties to characterize stability boundaries without requiring the construction of appropriate Lyapunov functions. The key innovation in the presented approach is the use of the path integral formulation to compute the Koopman principal eigenfunction, which enables lower-dimensional analysis without requiring calculation of the higher-dimensional Koopman operator. Unlike the SOS three-step iterative method for controlled invariant set computation, this approach requires only the computation of a single eigenfunction associated with the unstable eigenvalue of a type-1 unstable equilibrium point. This offers a potentially more direct path to stability boundary computation without the need for Lyapunov function verification and refinement. To verify the validity of the results from this new approach, a parametric study is conducted, which reveals how the stability region changes with different vehicle characteristics. The findings show that increased longitudinal velocity leads to a reshaping of the stability boundary in the $v_y - r$ phase plane, allowing for higher lateral velocities but restricting the maximum safe yaw rate. This observation aligns with expectations: at higher speeds, vehicles can achieve greater lateral motion but become more sensitive to angular velocity changes, requiring tighter constraints on yaw rate to maintain stability [18].

A notable advantage of the Koopman-based approach is its ability to address the temporal limitations present in the Lyapunov-based method. While Lyapunov stability theory concerns itself with asymptotic convergence over an infinite time horizon, the Koopman spectral approach can be applied to finite-time stability analysis. The method presented allows for the characterization of stability regions within specific time windows, which is particularly relevant for vehicle safety applications where the time scale of recovering from perturbations is critical. This finite-time perspective provides a more practical framework for emergency scenarios where rapid recovery is essential, addressing a key limitation in Lyapunov approaches, which offer no guarantees on recovery time despite ensuring eventual convergence.

The practical applicability of reachability analysis for online verification in autonomous vehicles has been demonstrated [8]. This implementation on a Cadillac SRX test vehicle showcases the feasibility of performing formal verification during actual vehicle operation. By computing reachable sets for both the ego vehicle and other traffic participants, their method formally verifies safety with respect to road boundaries and potential collisions. The approach presented differs from the database method by performing direct online computation of reachable sets using conservative linearization techniques, achieving computation speeds 1.79 times faster than real-time execution of maneuvers.

However, the online approach exhibits several significant limitations. To achieve real-time computation, a conservative linearization of the nonlinear vehicle dynamics is employed, which inevitably introduces overapproximation in the reachable sets. The method linearizes the nonlinear system around operating points and compensates for linearization errors through additional uncertainty terms. While this strategy enables faster computation, it produces overly conservative results, particularly in scenarios demanding aggressive maneuvers where nonlinearities become more pronounced. Their validation experiments are constrained to a narrow set of scenarios, with the double-lane change maneuver serving as the primary test case. Although testing confirmed controller functionality, these tests operated within a limited envelope. According to their results, the vehicle never experienced slip angles exceeding 6 degrees during testing. While such slip angles are expected during normal driving conditions, a safety controller must perform effectively outside the normal driving regime — precisely the circumstances in which it would be activated.

The field of lateral stability analysis for vehicles has progressed considerably, moving from simple linear models toward sophisticated nonlinear approaches that better capture real-world vehicle dynamics. Phase plane methods—particularly those using sideslip angle representations—have emerged as valuable tools for understanding stability boundaries. While Lyapunov-based methods offer mathematical guarantees of asymptotic stability, they lack temporal performance considerations crucial for emergency scenarios. Koopman operator approaches address this limitation by enabling finite-time stability analysis, providing more practical insights for safety-critical applications. Meanwhile, online reachability analysis demonstrates promising real-time verification capabilities, though current implementations remain conservative due to linearization compromises. However, the HJ reachability approach remains largely unused. As Section 1.1 explains, there is a lot of potential in using HJ reachability methods; therefore, the application to vehicle lateral dynamics warrants further study.

# 2

# Project Plan

## 2.1. Research Gap and Proposed Approach

After reviewing the literature, a research gap has been identified in the application of Hamilton-Jacobi (HJ) reachability analysis to vehicle lateral stability problems. While HJ reachability has been successfully applied in various domains, including aircraft safety systems, its potential for enhancing vehicle lateral stability has remained largely unexplored.

### 2.1.1. Research Gap

The review of existing literature reveals several key limitations in current approaches:

1. **Phase Plane Methods**: While phase plane methods provide valuable visual interpretations of stability regions, most lack formal mathematical guarantees and rely heavily on visual interpretation rather than rigorous computation.
2. **Lyapunov-Based Approaches**: Lyapunov-based approaches do provide mathematical guarantees for stability. However, the guarantees are for infinite time horizons, meaning they do not necessarily guarantee safety. Additionally, finding the Lyapunov functions becomes difficult with higher-dimensional systems.
3. **Koopman Operator Approaches**: Stability boundary calculations based on the spectral analysis of the Koopman Operator do provide mathematical guarantees for stability, on limited time horizons. However, working with the Koopman Operator is not straightforward, especially for higher-dimensional systems. Additionally, there are multiple teams working on developing these techniques for vehicle stability, so the research gap is being filled.

With a single exception, which focused on a linearized system and was meant to work in real time, Hamilton-Jacobi reachability has not been applied to vehicle lateral stability. HJ reachability is specifically meant for non-linear system analysis, and with recent developments in the field, higher dimension system analysis is becoming feasible [19].

### 2.1.2. Proposed Approach

Based on the identified research gap, this thesis aims to develop a framework for computing, storing, and utilizing Hamilton-Jacobi reachability-based safety envelopes for vehicle lateral stability. The proposed approach consists of several key components:

1. **Hamilton-Jacobi Reachability for Safe Sets**: Compute the safe sets of a nonlinear bicycle model for different velocity and road configurations using HJ reachability analysis. This should overcome limitations of previous methods by giving safe sets for a limited time frame, and can, in theory, include more than just the lateral dynamics (for example, actuator dynamics or longitudinal dynamics).
2. **Interpolation Between Safe Sets**: Implement and validate methods to interpolate between pre-computed safe sets for different operating conditions (velocities, road friction coefficients).
3. **Safety Controller Design**: Create a safety controller that utilizes these interpolated safe sets to ensure vehicle stability across varying conditions. While maintaining the safety guarantees for the interpolated safe sets is beyond the scope of the project, the created controller should demonstrate effective action under a variety of conditions and maneuvers.

4. **Validation on high-fidelity Model**: To check if the created controller is, in principle, transferable to real-life applications, a test using a higher fidelity simulation will be performed (CarMaker specifically). This should provide insight into whether simplified models are sufficient for creating a controller, or whether the underlying HJ reachability analysis needs to be closer to reality, for it to be useful.

## 2.2. Research Question(s)

Based on the identified research gap in vehicle stability control using Hamilton-Jacobi reachability analysis, this project addresses the following main research question:

*How can Hamilton-Jacobi reachability analysis be effectively applied to compute safe operating envelopes of a non-linear vehicle bicycle model and subsequently used in creating an effective safety controller?*

To systematically address this comprehensive question, the following sub-questions have been formulated:

1. **Computation of Safe Sets:** How can Hamilton-Jacobi reachability analysis be applied to compute accurate safe operating envelopes for the nonlinear bicycle model under different velocity and road friction conditions?
    - What time horizon provides the most practical balance between safety guarantees and operational freedom?
    - How do vehicle parameters and road conditions affect the shape and size of the computed safe sets?
    - What state representation (phase plane) is most appropriate for computing these safe sets?
2. **Interpolation Between Safe Sets:** What methods can be used to accurately interpolate between precomputed safe sets?
    - How can we ensure that interpolated sets preserve the controlled invariance property of the original sets?
    - What interpolation errors arise between different operating conditions, and how can these errors be minimized?
    - Can formal guarantees be established for the interpolated safe sets, and what are their limitations?
3. **Safety Controller Design:** How can the computed and interpolated safe sets be effectively utilized in a safety controller for vehicle lateral stability?
    - What control architecture best leverages the safe set information for real-time intervention?
    - How can smooth transitions be ensured when the operating conditions change and different safe sets become relevant?
    - What performance metrics should be used to evaluate the effectiveness of the safety controller?
4. **Validation of Safety Controller:** How can the developed safety controller be checked to be applicable in real-world scenarios?
    - To what extent does the safety controller ensure safety when used on a high-fidelity simulation?
    - How does the safety controller compare to other control methods when used in a high-fidelity simulation?
    - What steps can be taken to improve the safety and/or performance of the designed controller?

## 2.3. Methods, Tools, and Expected Results

This research will employ a systematic approach combining theoretical analysis, numerical computation, and simulation-based validation to address the research questions. The following sections detail the methodological framework, computational setup, and expected outcomes of the project.

### 2.3.1. Theoretical Framework

The theoretical foundation of this work rests on two primary mathematical frameworks: Hamilton-Jacobi (HJ) reachability analysis and nonlinear vehicle dynamics.

**Hamilton-Jacobi Reachability Analysis**
HJ reachability analysis provides a rigorous approach for computing controlled invariant sets for nonlinear systems. The key components of this framework include:

- **Backward Reachable Sets:** These sets represent states from which a target set (such as a stable equilibrium region) can be reached within a specified time horizon while satisfying system dynamics and constraints.
- **Forward Reachable Sets:** These sets represent states that can be reached from a given initial set within a specified time horizon while satisfying system dynamics and constraints.
- **Level Set Method:** This numerical method will be used to solve the Hamilton-Jacobi-Isaacs partial differential equation that defines the reachable sets. The method works by representing the boundary of the reachable set as the zero level set of an implicit surface function.

The safe flight envelope, as computed using DEFEND, is defined as the intersection of backward and forward reachable sets, representing states from which the system can both reach the target equilibrium and maintain maneuverability within a given time horizon [11]. This definition holds for time-reversible systems. While this is not necessarily the case for the vehicle models used, the FRS and BRS intersection should provide a good, but overestimated, safe set.

**Vehicle Dynamics Modeling**
For the vehicle dynamics, this research will employ:

- **Nonlinear Bicycle Model:** Nonlinear bicycle model has been established as the way to conduct lateral stabilization. During the thesis, the plan is to slowly build up the complexity of the vehicle model. At the start, the model will include a simple tire model (Dugoff, for example), no changes in longitudinal velocity, and no actuator dynamics. And as the thesis work progresses, more sophisticated tire models can be implemented, interaction between lateral and longitudinal velocities, and actuator dynamics can be included. The final form of the model used will depend on the speed of progress, but this iterative approach should allow for continuous progress and might give insight into which aspects of the vehicle model are most important for predicting safety

### 2.3.2. Computational Methods and Tools

**Software Tools**
The research will utilize the following software tools:

- **MATLAB/Simulink:** Primary environment for implementing numerical algorithms, vehicle models, and the control framework.
- **Level Set Toolbox with helperOC:** The Level Set Toolbox is an open-source MATLAB toolbox developed for solving Hamilton-Jacobi equations and computing reachable sets [20]. helperOC is a wrapper that simplifies the implementation [2].
- **CarMaker:** High-fidelity vehicle dynamics simulation software that will be used to validate the safety controller's performance in realistic scenarios.
- **Python:** For supplementary data analysis, visualization, and prototyping of interpolation algorithms. It is also possible that Python will be used to compute the safety sets, depending on what will be most appropriate.

**Computational Setup**
The computational framework will be structured as follows:

- **Offline Computation:** Safe sets will be computed offline for a discrete set of operating conditions (velocities and road friction coefficients) with the nonlinear bicycle model.
- **Database Construction:** Following the DEFEND approach, the computed safe sets will be stored in a structured database for efficient retrieval during simulation and eventual real-time implementation [11].
- **Interpolation Algorithms:** Custom algorithms will be developed to interpolate between the precomputed safe sets, with formal analysis of the interpolation error and preservation of safety properties.
- **Controller Implementation:** The safety controller will be implemented in Simulink and integrated with the CarMaker simulation environment for testing.

**Validation Methodology**
The validation of the approach will follow a three-stage process:

- **Numerical Simulation:** Testing of the approach using the nonlinear bicycle model to verify the computational accuracy and controller performance.
- **High-Fidelity Simulation:** Testing in CarMaker with various maneuvers (e.g., double lane change, slalom) under different road conditions to evaluate real-world performance.

### 2.3.3. Expected Results and Data
The primary outcomes of this research are expected to be:

1. A comprehensive set of HJ reachability-based safe operating envelopes for vehicle lateral stability under various conditions.
2. An algorithm for interpolating between safe sets while preserving safety guarantees, with quantifiable error bounds.
3. A safety controller architecture that utilizes the precomputed and interpolated safe sets to ensure vehicle stability.
4. Validation of the controller on a high-fidelity simulation.

These outcomes will contribute to both the theoretical understanding of vehicle stability control and practical implementations for advanced driver assistance systems and autonomous vehicles.

## 2.4. Planning
This section outlines the project timeline, division of work packages, milestones, and contingency plans to ensure the successful completion of the research. The project has been organized into logical work packages with clear dependencies, estimated durations, and associated deliverables.

### 2.4.1. Work Packages
The project is divided into the following work packages (WP):

WP1: **Literature Review and Theoretical Framework Refinement** (6 weeks)

- A review of safety filters and vehicle stability approaches
- Preparation of a plan for the Master's thesis
- Deliverable: A project plan, including an overview of the literature

WP2: **HJ Reachability Calculation Set-Up** (2 weeks)

- Create the basic nonlinear bicycle model
- Compute first forwards and backwards reachable sets
- Deliverable: Preliminary safety sets and a code base ready for further work

WP3: **Safe Set Database Generation** (4 weeks)

- Computation of safe sets for various operating conditions
- Systematic organization of computed sets in a structured database
- Analysis of safe set characteristics across different conditions
- Deliverable: Database of precomputed safe sets

WP4: **Interpolation Method Development** (8 weeks)

- Design of interpolation algorithms for safe sets
- Implementation and testing of interpolation methods
- Formal analysis of interpolation error and safety preservation
- Deliverable: Interpolation algorithm with validation report

WP5: **Safety Controller Design** (6 weeks)

- Development of control architecture utilizing safe sets
- Implementation of blending mechanisms for smooth intervention
- Integration with the database and interpolation system
- Deliverable: Controller implementation

WP6: **Controller Validation in Simulation** (4 weeks)

- Integration with CarMaker high-fidelity simulation environment
- Testing across various scenarios and operating conditions
- Quantitative performance evaluation and comparison
- Deliverable: Validation report with performance metrics

WP7: **Analysis and Refinement** (4 weeks)

- Comprehensive analysis of results
- Refinement of methods based on validation outcomes
- Identification of limitations and future work directions
- Deliverable: Analysis report with recommendations

WP8: **Thesis Writing and Finalization** (6 weeks)

- Compilation of all research findings
- Drafting and revision of thesis chapters
- Preparation of presentations and demonstration materials
- Deliverable: Complete thesis manuscript and presentation materials

### 2.4.2. Milestones and Gantt Chart

Following the TU Delft Aerospace Engineering Master Thesis Timeline, the project is structured around the following formal milestone meetings and deliverables:

- **Kick-off Meeting**
  - Clarification of research topic and scope
  - Agreement on expectations, timeframe, and supervision arrangements
  - Discussion of required resources and potential constraints
- **Research Proposal Review**
  - Submission of Research Proposal deliverable
  - Evaluation of literature review and state-of-the-art understanding
  - Assessment of SMART research questions and hypotheses
  - Review of proposed research activities and project plan
- **Mid-term Review**
  - Submission of Mid-term Deliverable
  - Evaluation of initial research results from Phase 1
  - Critical analysis of results in relation to research questions
  - Assessment of progress against original schedule
  - Review of thesis outline and completed writing
- **Green Light Review**
  - Submission of Draft Thesis
  - Comprehensive evaluation of thesis completeness
  - Assessment of research documentation, results presentation, and critical discussion
  - Go/no-go decision for thesis defense

- **Request Examination and Thesis Hand-In**
  - Formal request for thesis examination
  - Scheduling of defense date
  - Submission of all research data, models, and supplementary documentation
- **Thesis Defense**
  - Final presentation and defense of research
  - Completion of all degree requirements

All the work packages and milestones are presented in Figure 2.1.

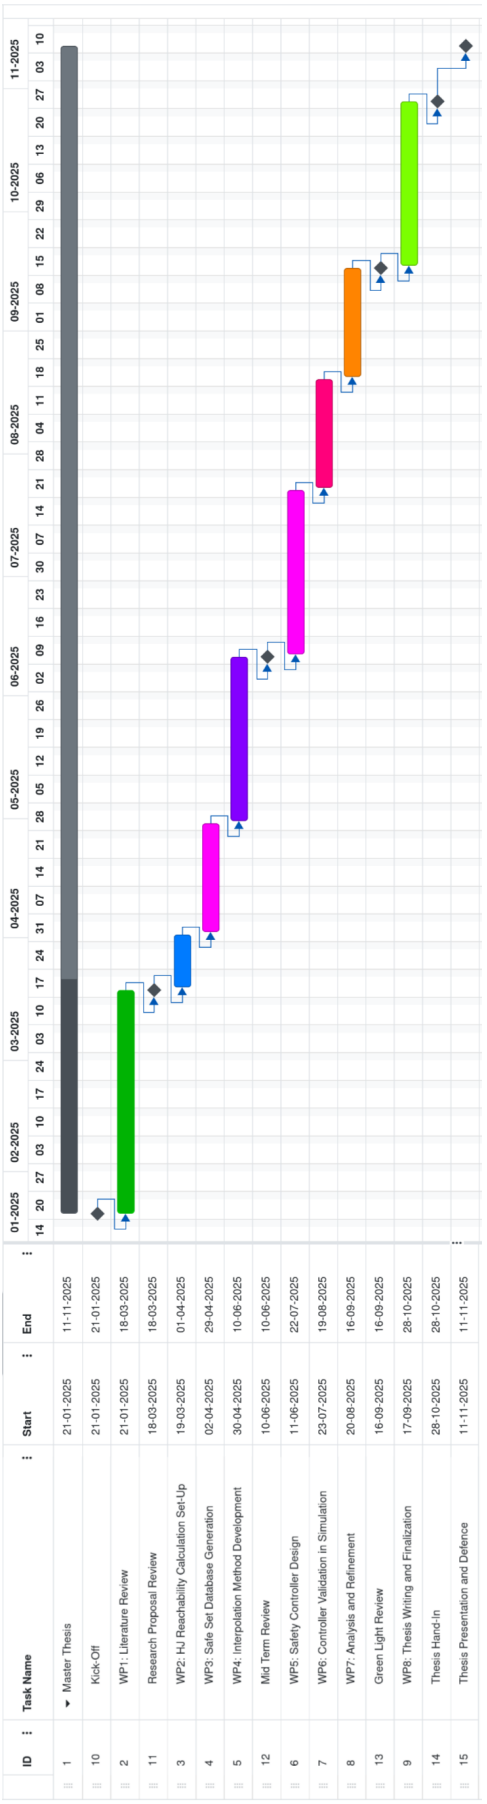


**Figure 2.1:** Project timeline Gantt Chart

# Part I

## Scientific Article

# Hamilton-Jacobi Reachability for Vehicle Lateral Stability: A Control Barrier-Value Function Approach

B. Jemioł *
*TU Delft, Delft, Netherlands*

**Traditional electronic stability control systems exhibit fundamental limitations when vehicles operate at their handling limits due to nonlinear tire saturation effects and reactive threshold-based interventions that can further destabilize vehicles during emergency maneuvers. This research develops and validates a proactive vehicle lateral stability controller using Control Barrier-Value Function Quadratic Programming that combines Hamilton-Jacobi reachability analysis with control barrier function frameworks. The approach computes controlled invariant sets through backward reachable tube analysis using a nonlinear bicycle model. These sets are then used to synthesize safety constraints that maintain vehicle stability without sacrificing performance. The resulting controller operates by solving a quadratic program that minimally modifies nominal control inputs only when necessary to maintain safety margins. Validation through high-fidelity CarMaker simulation using a passenger vehicle model demonstrates the controller's effectiveness across varying road conditions, including high-grip surfaces ($\mu$ = 1.0), low-grip surfaces ($\mu$ = 0.2), and dynamic friction transitions from high to low grip. The CBVF safety filter achieved up to 82% improvement in tracking performance and 75% reduction in control effort compared to baseline LQR control, while successfully preventing vehicle spin-out scenarios that traditional controllers failed. Critically, the controller maintained stable performance across all tested conditions despite being tuned only for high-grip scenarios, demonstrating significant parameter error resilience and eliminating the need for surface-specific calibration. The parameter study revealed that different combinations of maximum control authority, prediction horizon, and relaxation parameters yield distinct control behaviors ranging from boundary-focused safety enforcement to steady-state tracking enhancement. These results establish Control Barrier-Value Function Quadratic Programming as a viable augmentation framework for existing vehicle stability control architectures, offering a practical approach for safety-critical automotive applications.**

## I. Introduction

The development of electronic stability control systems has been the biggest contributing factor to reducing road fatalities, since the adoption of seat belts and airbags [1, 2]. Mercedes-Benz introduced the first production ESC system in 1995, establishing differential braking as the primary intervention mechanism [3]. Despite widespread adoption, ESC became mandatory for all U.S. vehicles in 2012, limitations persist when vehicles operate at their handling limits [4, 5]. The limitations of linear control methods become apparent during emergency maneuvers where tire forces saturate nonlinearly [5]. Nonlinear tire saturation effects fundamentally alter vehicle stability characteristics when lateral forces approach the friction circle limits. At low slip angles (< 0.1 rad), tire lateral force exhibits linear behavior characterized by cornering stiffness coefficients [6]. However, as slip angles increase beyond 0.15-0.2 rad, tire forces reach peak values and subsequently decrease, creating a nonmonotonic relationship between slip angle and lateral force.

Traditional ESC systems activate based on predetermined thresholds of yaw rate error or sideslip angle, resulting in abrupt interventions that can destabilize the vehicle further [7]. This reactive approach fails to account for the complex, state-dependent nature of vehicle dynamics at high lateral accelerations. The phase-plane analyses developed in the 1970s provided visualization of stability boundaries but offered limited practical guidance for controller design in multi-dimensional state spaces [8, 9]. To address these fundamental limitations of threshold-based reactive control, researchers have turned to controlled invariant set theory as a proactive framework for ensuring safety.

Controlled invariant sets are regions of the state space where appropriate control inputs can maintain vehicle stability indefinitely. The first attempts at computing controlled invariant sets for vehicles used sum-of-squares (SOS) programming to compute maximum controlled invariant sets for nonlinear vehicle models, based on Lyapunov stability [10]. The approach successfully demonstrated that different actuators (steering versus differential braking)

*Master's Student, Control & Simulation, Faculty of Aerospace Engineering

produce distinct stability regions, providing theoretical justification for coordinated control strategies. However, this approach has limitations. First, the controlled invariant sets are computed over infinite time horizons, meaning the method includes states from which recovery may require impractically long durations. A vehicle state might technically belong to the invariant set yet require 30 seconds of maximum steering input to stabilize. Second, their problem formulation couples steering angle and steering rate constraints, producing the counterintuitive result that increasing actuator rate limits reduces the size of the controlled invariant set. This artifact of the mathematical formulation, rather than physical reality, undermines applicability.

The infinite-horizon limitation has been addressed through Koopman spectral analysis, enabling computation of time-limited controlled invariant sets [11]. The Koopman approach requires computation of only one eigenfunction associated with the unstable eigenvalue of the type-1 unstable equilibrium point, offering computational advantages over traditional Lyapunov methods. By leveraging path integral formulations, they construct stability boundaries that explicitly account for finite-time recovery constraints. Time-limited invariant sets directly address the practical question: from which states can the vehicle be stabilized within a specified duration? The results demonstrate that stability regions shrink significantly when recovery time constraints are imposed, highlighting the optimism bias in traditional infinite-horizon analyses. Yet the Koopman approach inherits the challenges of Lyapunov methods. Constructing appropriate basis functions that accurately capture nonlinear dynamics remains difficult.

Another approach for finding controlled invariant sets is Hamilton-Jacobi reachability analysis. Unlike Lyapunov methods requiring careful function construction, HJ analysis directly computes the value function representing the minimum time to reach unsafe states [12]. The resulting level sets provide controlled invariant boundaries without manual tuning. HJ reachability has been demonstrated for vehicle safety verification, computing reachable sets online to ensure safety [13]. Their experimental validation on a test vehicle proved feasibility. Though the method was limited to the use of linearized models to enable online operation. Recent developments in HJ computation, including local update schemes and GPU acceleration, have reduced computational barriers that previously limited implementation [14].

Parallel to HJ developments, Control Barrier Functions (CBFs) emerged as an alternative framework encoding safety through controlled invariance constraints [15]. Extending barrier certificates from verification to control synthesis, CBFs provide a systematic method for modifying control inputs to ensure safety [16]. The key insight was recognizing that for control-affine systems, CBF constraints become linear in the control input, enabling real-time safety filtering through quadratic programs (QP) [17]. Unlike HJ methods producing bang-bang control, CBF-QPs generate smooth control signals that respect actuator limitations [18]. Adaptive cruise control and lane keeping have been demonstrated using CBFs, showing how multiple safety constraints could be simultaneously enforced through a single optimization [15]. The ability to minimally modify a nominal controller, intervening only when necessary, proved particularly attractive for integrating safety into existing control architectures [19]. However, CBF methods face similar challenges to Lyapunov and Koopman approaches. Constructing valid barrier functions for general nonlinear systems is difficult. While simple constraints (maintaining lane boundaries, following distance) admit straightforward CBF formulations, complex stability constraints require careful design [18]. The conservatism inherent in hand-crafted barrier functions often results in unnecessarily restricted operating regions [4]. Moreover, verifying that a candidate function satisfies CBF conditions across the entire state space remains computationally intensive for high-dimensional systems.

Hamilton-Jacobi reachability and Control Barrier Functions (CBFs) represent the two dominant safety filter frameworks, each with complementary strengths [20]. HJ reachability provides systematic computation of safe sets but produces conservative controllers with jerky inputs. CBF controllers offer smooth interventions through quadratic program formulations but require manual barrier function design, often resulting in overly conservative safe regions. Recent work demonstrates these approaches can be unified with HJ value functions that satisfy CBF conditions, enabling systematic construction of smooth safety controllers [21]. Combining the strengths of both approaches, while maintaining formal safety guarantees, promises to be the ultimate in safety controllers.

The Control Barrier-Value Function Quadratic Programming controller (CBVF-QP) represents a significant advancement in this unification effort [21]. This approach merges HJ-derived value functions with CBF constraint structures, yielding a safety controller that maintains finite-time safety guarantees while providing smooth control actions. The CBVF-QP inherits the systematic constructive properties of HJ reachability, eliminating the hand-crafting requirements of traditional CBF approaches. This methodology produces the maximal safe set for desired safety constraints while ensuring control bounds are satisfied throughout the safe region, offering both theoretical rigor and practical implementability for safety-critical control applications.

This paper advances vehicle lateral stability control through systematic application of HJ reachability analysis and creation and testing of a CBVF-QP safety controller. The contributions of this paper are as follows: (1) compute controlled invariant sets using HJ reachability, capturing nonlinear tire saturation effects; (2) develop a controller using HJ-derived CBF constraints, ensuring persistent safety without sacrificing performance; and (3) validate the controller through high-fidelity simulation, quantifying performance in both safety margins and control smoothness.

## II. Background

This section provides the theoretical foundations for computing safety envelopes using Hamilton-Jacobi (HJ) reachability analysis, Control Barrier Function, and how they can be combined into a safety controller architecture.

### A. Hamilton-Jacobi Reachability

Hamilton-Jacobi reachability analysis provides a mathematical framework for computing controlled reachable sets for nonlinear dynamical systems [12]. The goal of reachability analysis is to establish which states can be reached, avoided, or recovered from, given a dynamical system. For this research, the main area of interest lies in the Backwards Reachable Tube. In general, BRTs are sets of states for which a dynamic system, under a given control law and time constraint, can achieve a set of target states. For example, a BRT of a Dubins car robot, given a small area as a target, will give you all the positions from which the robot can start, and reach that target within the specified time. The formal mathematical approach is given below.

Consider a general nonlinear system as in Equation 1,

$$\dot{x} = f(x, u, d), \quad x \in \mathbb{R}^n, \quad u \in \mathcal{U}, \quad d \in \mathcal{D} \tag{1}$$

where $x$ represents the system state, $u$ denotes the control input constrained to set $\mathcal{U}$ and $d$ denotes the disturbance constrained to $\mathcal{D}$. The fundamental concept in HJ reachability is the computation of reachable tubes that characterize all possible system trajectories under optimal control and worst-case disturbances. The Backward Reachable Tube represents all initial states from which the system can be driven to a target set $\mathcal{L}$ within a time horizon $[0, T]$.

Formally, the BRT is defined by Equation 2,

$$\mathcal{R}(T, \mathcal{L}) = \{x_0 : \exists u(\cdot) \in \mathcal{U}, \forall d(\cdot) \in \mathcal{D}, \exists t \in [0, T], \xi(t) \in \mathcal{L}\} \tag{2}$$

where $\xi(t)$ denotes the system trajectory starting from $x_0$ under control $u(\cdot)$. The BRT can be computed by solving the Hamilton-Jacobi-Isaacs (HJI) partial differential equation, shown in Equation 3,

$$\frac{\partial V}{\partial t}(x, t) + \min\{0, H(x, \nabla V(x, t))\} = 0 \tag{3}$$

with terminal condition $V(x, 0) = l(x)$, where $l(x)$ is the implicit surface function representing the target set $\mathcal{L}$. The Hamiltonian $H$ is given by Equation 4.

$$H(x, p) = \min_{u \in \mathcal{U}} \max_{d \in \mathcal{D}} p \cdot f(x, u, d) \tag{4}$$

The zero level set of the value function $V(x, T)$ defines the boundary of the BRT: $\mathcal{R}(T, \mathcal{L}) = \{x : V(x, T) \leq 0\}$. The obtained value function $V(x, T)$ is the most important part of the reachability analysis for this application. As will be discussed shortly, it has all the properties that are needed for a Control Barrier Function, and having a methodical approach to obtaining the barrier function would eliminate one of the biggest challenges in CBF safety controller creation.

The complex nature of the HJI PDE means that finding an analytical solution is, in most cases, impossible. Therefore, numerical methods are employed to solve for the value function. The numerical solution of the HJI PDE employs the level set method, which represents the reachable set boundary as the zero level set of an implicit surface function. This approach handles topological changes in the reachable set and provides stable numerical schemes for nonlinear systems [22]. The method discretizes the state space on a grid and propagates the value function using upwind finite difference schemes that respect the direction of information flow determined by the system dynamics.

### B. Control Barrier Functions

Control Barrier Functions (CBFs) provide a powerful framework for synthesizing safety-critical controllers that ensure control forward invariance of sets for nonlinear systems [15, 23]. Unlike traditional approaches that rely on post-hoc verification, CBFs enable the design of controllers that inherently guarantee safety by construction.

Consider an affine nonlinear control system:

$$\dot{x} = f(x) + g(x)u, \quad x(0) = x_0 \in \mathcal{X}_0 \tag{5}$$

where $x \in \mathcal{X} \subset \mathbb{R}^n$ is the state, $u \in \mathbb{R}^m$ is the control input, and $f(x) : \mathbb{R}^n \to \mathbb{R}^n$ and $g(x) : \mathbb{R}^n \to \mathbb{R}^{n \times m}$ are assumed to be locally Lipschitz continuous. The state space $\mathcal{X}$ can be decomposed into a desired set $\mathcal{X}_d$ and undesired set $\mathcal{X}_u$, such that $\mathcal{X}_d \cup \mathcal{X}_u = \mathcal{X}$.

A continuously differentiable function $h : \mathbb{R}^n \to \mathbb{R}$ characterizes the safe set such that:

$$\begin{aligned} \mathcal{X}_d &= \{x \in \mathcal{X} : h(x) \geq 0\} \\ \partial\mathcal{X}_d &= \{x \in \mathcal{X} : h(x) = 0\} \\ \text{Int}(\mathcal{X}_d) &= \{x \in \mathcal{X} : h(x) > 0\} \end{aligned} \tag{6}$$

where $\partial\mathcal{X}_d$ and $\text{Int}(\mathcal{X}_d)$ denote the boundary and interior of the invariant set, respectively. The continuously differentiable function $h : \mathbb{R}^n \to \mathbb{R}$ is a Control Barrier Function (CBF) for a relative degree one system if there exists an extended class $\mathcal{K}_e$ function $\alpha$ such that for all $x \in \mathcal{X}$, Equation 7 holds. An extended class $\mathcal{K}_e$ function $\alpha : \mathbb{R} \to \mathbb{R}$ is a strictly increasing function satisfying $\alpha(0) = 0$, which extends the standard class $\mathcal{K}$ definition to the entire real line.

$$\sup_{u \in \mathcal{U}} [L_f h(x) + L_g h(x) u + \alpha(h(x))] \geq 0 \tag{7}$$

In Equation 7, $L_f h(x) = \frac{\partial h}{\partial x} f(x)$ and $L_g h(x) = \frac{\partial h}{\partial x} g(x)$ are the Lie derivatives of $h$ along the vector fields $f$ and $g$, respectively. This formulation applies when $L_g h(x) \neq 0$, ensuring the control input appears explicitly in $\dot{h}$. For systems with a higher relative degree, the CBF condition must be extended to higher-order derivatives. The key insight of CBF-based control is that any Lipschitz continuous controller $u^*(x)$ satisfying the CBF constraint renders the desired set $\mathcal{X}_d$ forward invariant. The forward invariance property emerges from analyzing the time derivative of $h$ along system trajectories. For any trajectory $x(t)$ of the closed-loop system with a controller $u^*$ satisfying the CBF constraint, we have:

$$\dot{h}(x(t)) = \frac{\partial h}{\partial x} \dot{x} = L_f h(x) + L_g h(x) u^* \geq -\alpha(h(x)) \tag{8}$$

This inequality ensures that whenever the system state approaches the boundary of the desired set (i.e., as $h(x) \to 0^+$), the derivative $\dot{h}(x)$ remains non-negative since $\alpha(0) = 0$ for extended class $\mathcal{K}_e$ functions. Consequently, the value of $h$ cannot decrease below zero, preventing trajectories from leaving the invariant set. This result can be viewed as an application of Nagumo's theorem, which characterizes forward invariance through constraints on the vector field at the boundary of the set.

This forward invariance constraint enables the integration of safety constraints directly into optimization-based control frameworks, most commonly through Quadratic Programming (QP):

$$u^* = \arg\min_{u \in \mathcal{U}} \quad \frac{1}{2} \|u - u_{\text{ref}}\|^2 \tag{9}$$

$$\text{s.t.} \quad L_f h(x) + L_g h(x) u + \alpha(h(x)) \geq 0 \tag{10}$$

where $u_{\text{ref}}$ is a reference controller designed for performance objectives. Formulating the safety controller in such a manner has several key advantages. The optimization ensures minimal intervention, meaning the controller only modifies the nominal control when necessary for safety, preserving performance objectives when possible. The QP problem can be solved in real-time using standard convex optimization solvers, making the controller computationally efficient. Lastly, the controller is modular, since the safety layer (CBF constraint) is decoupled from the performance controller design.

### C. HJ Value Function as a Control Barrier Function

The Control Barrier-Value Function (CBVF) framework was developed to unify HJ reachability analysis and CBF methods [21]. This framework demonstrates that HJ value functions can serve directly as barrier functions while enabling less conservative control synthesis than standard reachability approaches. The key innovation is the introduction of an exponential discount factor $\gamma \geq 0$ into the standard HJ reachability formulation [21]. The CBVF is defined as:

$$B_\gamma(x, t) = \min_{d \in \mathcal{D}} \max_{u \in \mathcal{U}} \min_{s \in [t, 0]} e^{\gamma(s-t)} l(x(s)) \tag{11}$$

where $l(x)$ represents the constraint function defining the safe set. When $\gamma = 0$, this reduces to the standard HJ value function. For $\gamma > 0$, the exponential term modifies the optimal control policy in a manner analogous to—but distinct from—exponential CBFs.

It has been proven that $B_\gamma$ is the viscosity solution to the following Hamilton-Jacobi-Isaacs Variational Inequality (HJI-VI), where $D_t$ denotes the time derivative operator and $D_x$ denotes the spatial gradient operator [21]:

$$0 = \min \left\{ l(x) - B_\gamma(x, t), \quad D_t B_\gamma(x, t) + \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} D_x B_\gamma(x, t) \cdot f(x, u, d) + \gamma B_\gamma(x, t) \right\} \tag{12}$$

The mathematical guarantees demonstrate that this framework preserves the safety properties of HJ reachability [21]. Specifically, the zero-superlevel set $C_\gamma(t) := \{x \in \mathbb{R}^n : B_\gamma(x,t) \geq 0\}$ equals the viability kernel $\mathcal{S}(t)$ for each time $t \leq 0$. This ensures that the CBVF approach recovers the maximal safe set, addressing the primary limitation of hand-crafted CBFs that often yield conservative safe regions.

The critical advantage of introducing $\gamma$ emerges in the resulting control constraints. Standard HJ reachability optimal control requires the constraint function to remain non-decreasing along trajectories: $\dot{l}(x(t)) \geq 0$. This restriction forces the system to always move away from the safety boundary, resulting in overly conservative behavior. In contrast, the CBVF framework permits controlled decay: $\dot{l}(x(t)) \geq -\gamma l(x(t))$. This relaxation allows the system to approach the safety boundary at a controlled rate while maintaining safety guarantees, providing significantly more control freedom within the safe set.

For practical implementation on control-affine systems of the form $\dot{x} = p(x) + q(x)u + r(x)d$, a quadratic program formulation enables online control synthesis [21]:

$$\begin{aligned} u^* = \arg\min_{u \in \mathcal{U}} \quad & (u - u_{\text{ref}})^T (u - u_{\text{ref}}) \\ \text{s.t.} \quad & D_t B_\gamma(x,t) + D_x B_\gamma(x,t) \cdot [p(x) + q(x)u] + \min_{d \in \mathcal{D}} D_x B_\gamma(x,t) \cdot r(x)d + \gamma B_\gamma(x,t) \geq 0 \end{aligned} \tag{13}$$

This CBVF-QP controller has strong feasibility guarantees: it is feasible everywhere in the safe set where the gradient exists [21]. The resulting controller combines the real-time applicability of CBF-QP methods with the systematic barrier function construction of HJ reachability, while producing smooth control signals rather than the discontinuous bang-bang behavior typical of least-restrictive HJ controllers.

## III. Methodology

The goal of this research is to investigate whether CBVF can be used as an effective stability control system in vehicles. Ideally, the highest quality vehicle models should be used during BRT calculations. However, due to the "curse of dimensionality," finding the reachable set for systems above 10 degrees of freedom is computationally impractical [24]. This presents a problem, since vehicle models used for validating vehicle lateral stability controllers start at 14 degrees of freedom and can extend significantly beyond that [4, 25].

Therefore, the CBVF safety controller needs to be developed on a lower-order model. For modeling lateral dynamics, the bicycle model is the accepted reduced form of the problem. While there has not been work showing the validity of the reduction using traditional order reduction techniques, there is a vast body of experimental work showing that the bicycle model provides a good approximation of real-world behavior [7].

Recent developments in the HJ reachability space have shown that under assumptions of regularity, stability of fast dynamics, and Isaac's condition for the full model, reduced models can provide accurate reachable sets [26]. High-fidelity vehicle models do not necessarily fulfill all these assumptions (for example, nonlinear tire saturation violates the regularity assumption); therefore, no guarantee can be made about the performance of the safety controller. Nevertheless, this work demonstrates that a representative reduced model can perform well. Since the bicycle model has demonstrated its capability in lateral augmentation in vehicles, it can be predicted that a CBVF safety controller would work to some extent. Testing the CBVF safety filter on a high-fidelity model will provide insights into how much these simplifications affect the applicability of this controller synthesis.

### A. Bridging the Model Complexity Gap

The high-fidelity, high-order simulation is performed using IPG CarMaker with a validated passenger vehicle model [27]. While the exact formulation of the dynamic system is proprietary, the use of a validated vehicle model ensures that the results are representative of real-world behavior. The IPG CarMaker simulation provides a sufficient order gap between the 2-DOF bicycle model used for HJ reachability computations and a realistic full-order vehicle system, enabling empirical validation of the reduced-order safety controller approach.

The central challenge of this work is determining whether a safety controller synthesized from a reduced-order model can provide an effective safety margin when applied to higher-order systems. This question represents a fundamental tension in safety-critical control: while formal guarantees require exact models, computational tractability and the inherent differences between any model and reality necessitate practical compromises. Traditional safety controller frameworks rely on using identical models for both synthesis and application to maintain formal guarantees [18]. Using a reduced system for controller synthesis inherently relaxes these guarantees, which conflicts with the primary motivation for employing safety filters.

However, the persistent curse of dimensionality has motivated recent research toward reduced-order formulations with bounded error guarantees. Singular perturbation theory establishes bandwidth separation between fast and slow dynamics. Under assumptions of stable fast dynamics, the safety controller needs only to regulate the slower

dynamics to maintain system safety [26]. Koopman operator theory enables the transformation of nonlinear dynamics into a linear domain where reachability analysis becomes computationally tractable, though truncation of the infinite-dimensional representation introduces approximation errors [28]. Both approaches provide bounded error guarantees, but only for systems satisfying specific conditions: regularity, stable fast dynamics, and fulfillment of Isaac's condition.

High-fidelity vehicle models do not satisfy all these theoretical requirements. For instance, nonlinear tire saturation violates the regularity assumption, and the coupling between longitudinal and lateral dynamics complicates the separation of timescales. Consequently, no formal guarantees can be established for the bicycle-model-based safety controller when applied to the full-order simulation. This represents a knowledge gap in the current state of reduced-order safety control theory. Nevertheless, this is a rapidly evolving research area, and future work may extend theoretical guarantees to broader system classes.

Despite the lack of formal guarantees, the bicycle model has proven to be an effective tool for lateral stability augmentation through extensive experimental validation [7]. While formal order reduction analysis for vehicle dynamics remains an open research question, decades of practical application have confirmed that the bicycle model captures the safety-critical lateral dynamics with sufficient accuracy for control design. Extended models incorporating longitudinal velocity and weight transfer can improve performance, indicating that not all effects are captured, but the fundamental stability behavior is well represented.

The objective of this study is to empirically evaluate whether the bicycle model provides sufficient fidelity for CBVF-based safety filtering. By testing the controller on a high-fidelity simulation, this work establishes practical bounds on the applicability of reduced-order safety controllers and contributes empirical evidence to complement ongoing theoretical developments in this area.

### B. Nonlinear Vehicle Dynamics Modeling

Two models are used throughout this work. First, a nonlinear bicycle model with nonlinear tire dynamics for BRT computation and CBVF safety controller formulation. Second, for evaluating the performance of the safety controller, the high-fidelity CarMaker simulation described above is used. The vehicle parameters used throughout the work are shown in Table 1. The performance on the high-fidelity model should be indicative of potential real-world performance.

| Parameter | Value | Unit |
|---|---|---|
| $m$ | 1708 | kg |
| $V_x$ | 27.78 | m/s |
| $L_f$ | 1.536 | m |
| $L_r$ | 1.575 | m |
| $I_z$ | 2985.216 | kg·m$^2$ |
| $C_f$ | $1.5745 \times 10^5$ | N/rad |
| $C_r$ | $1.6426 \times 10^5$ | N/rad |

**Table 1 Vehicle parameters for BRT computation**

The nonlinear bicycle model considers two degrees of freedom: the yaw rate $\dot{\psi}$ and the vehicle sideslip angle $\beta$ [7]. The front steering angle $\delta$ is used as a parameter in the simulations. Since the controller will work through differential braking, the control input is a pure moment applied to the car. The nonlinear vehicle dynamics are expressed in state space form in Equation 14.

$$\dot{\mathbf{x}} = \begin{bmatrix} \ddot{\psi} \\ \dot{\beta} \\ \dot{\delta} \end{bmatrix} = \begin{bmatrix} \frac{1}{I_z}\left(W_f \sin\delta \cdot (F_{y,FL} - F_{y,FR}) + l_f \cos\delta \cdot (F_{y,FL} + F_{y,FR}) - l_r(F_{y,RL} + F_{y,RR})\right) \\ \frac{1}{mV_x}\left((F_{y,FL} + F_{y,FR}) \cdot \cos\delta + F_{y,RL} + F_{y,RR}\right) - \dot{\psi} \\ 0 \end{bmatrix} + \begin{bmatrix} \frac{1}{I_z} \\ 0 \\ 0 \end{bmatrix} u \quad (14)$$

Here, $I_z$ represents the vehicle yaw moment of inertia, $L_f$ and $L_r$ are the distances from the center of gravity to the front and rear axles respectively, $m$ is the vehicle mass, $W_f$ is the half front track width, and $V_x$ denotes the longitudinal velocity, assumed constant for lateral dynamics analysis. $F_{y,FL}$, $F_{y,FR}$, $F_{y,RL}$, $F_{y,RR}$ are the lateral tire forces for the front and rear, left and right sides.

The lateral tire forces are modeled using the magic formula, which captures the nonlinearities of tire behavior [29]. The general form is described in Equation 15.

$$F_y = f_{\text{magic}}(\alpha_f, F_{zf}, \mu, C_f) \tag{15}$$

where $F_{zf}$ and $F_{zr}$ represent the normal forces on the front and rear tires, $\mu$ is the tire-road friction coefficient, and $C_f$ and $C_r$ are the front and rear tire cornering stiffnesses.

The tire slip angles for the front and rear axles are computed using Equation 16.

$$\begin{aligned}
\alpha_{FL} &= -\delta + \arctan\left(\frac{\beta + l_f\dot{\psi}}{V_x - W_f\dot{\psi} + \varepsilon}\right) \\
\alpha_{FR} &= -\delta + \arctan\left(\frac{\beta + l_f\dot{\psi}}{V_x + W_f\dot{\psi} + \varepsilon}\right) \\
\alpha_{RL} &= \arctan\left(\frac{\beta - l_r\dot{\psi}}{V_x - W_r\dot{\psi} + \varepsilon}\right) \\
\alpha_{RR} &= \arctan\left(\frac{\beta - l_r\dot{\psi}}{V_x + W_r\dot{\psi} + \varepsilon}\right)
\end{aligned} \tag{16}$$

### C. Backwards Reachable Tube Computation

Since most HJ reachability problems do not have analytical solutions, numerical solvers are required. The chosen framework for computation is the `hj_reachability`* Python package developed by Stanford's *Autonomous Systems Lab*. The code needed to be modified to accommodate the relaxation parameter $\gamma$.

An important step in the BRT computation is the choice of the constraint function, also known as the target set. The constraint function defines the states from which the solution is computed and indicates the states that are required to be reachable, given the time and control limits. For this specific scenario, the target set represents the desirable states, and the BRT is the set of states that can reach anywhere in the constraint function, given the time limit and the limit of the applied moment. This choice of the constraint function is a design decision in the controller synthesis. In principle, any set of states can be chosen as the target set. However, for the CBVF-QP controller to satisfy its goals, the target should represent some sort of 'stable' or 'safe' region of the state space.

The choice of the origin would be natural, especially since most passenger vehicles are built with stability in mind (understeer behavior at the limit). However, this does not account for the effect of steerable front wheels. With front wheels deflected, it is expected that a vehicle will achieve some yaw rate and some sideslip. Therefore, the target should be a small tube around steady-state cornering points in the state space.

Acquiring steady-state cornering points can be done in many different ways, from geometric approaches like Ackermann steering to simulating the whole vehicle model until it achieves steady state. Given that the HJ reachability computation is expected to take significant computational resources, spending additional time on obtaining accurate steady-state solutions is not a concern. Therefore, the approach of simulating the model until it achieves a steady state was chosen.

The procedure to compute the steady state is as follows. The system is initiated with zero yaw rate and zero sideslip. The steering angle is slowly increased to the desired value and the system is allowed to settle in a stable condition, meaning the change in yaw rate and sideslip is lower than `1e-6` per time step. This is repeated for all discrete steering angles in the solution grid. The comparison of a zero-based target set to a steady-state set is shown in Figure 1.

The final part of computing the BRTs is parameter selection. The properties of the car come from the chosen model and are presented in Table 1. The maximal control moment, time horizon, and relaxation parameter $\gamma$ also need to be selected. During testing, these parameters proved to have an outsized effect on the results. These parameters have the potential to fundamentally change how the CBVF safety controller operates. This is why a parameter study was performed, and the results are presented in subsection IV.A.

### D. Baseline Controller

Since the CBVF safety controller is meant as an addition on top of existing solutions, a baseline controller was chosen. LQR is a popular, well-performing, and easy-to-implement controller [30, 31]. This controller generates a reference yaw rate to follow, based on the steering input and vehicle state, and uses variable-gain P control to track the reference. This formulation enables the controller to act as both a stability and maneuverability augmentation.

The reference generator is responsible for creating the yaw rate reference signal that the stability controller follows. The reference yaw rate is meant to align with driver expectations and also account for the nonlinear

*https://github.com/StanfordASL/hj_reachability

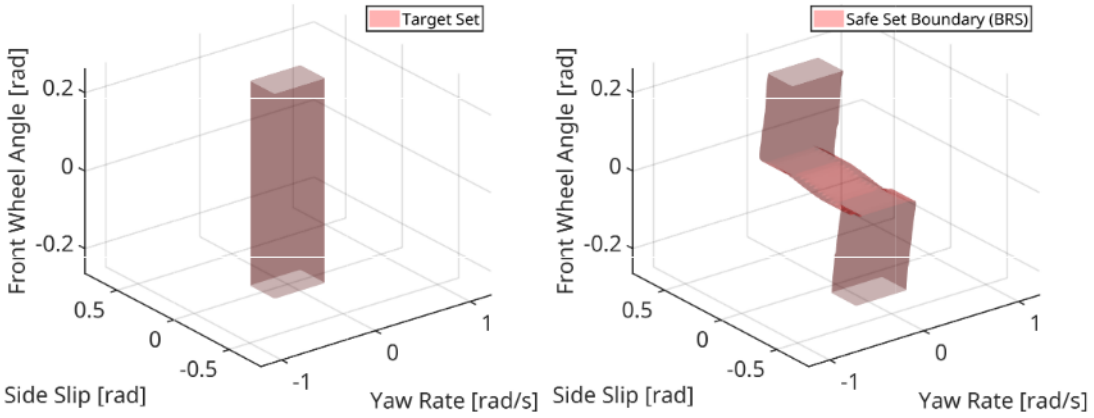


**(a) Without steady-state correction** **(b) With steady-state correction**

**Fig. 1 The constraint functions**

behavior of the vehicle. The generator operates in three stages: steady-state wheel angle to yaw rate approximation, then accounting for nonlinear tire saturation, and lastly approximating unmodeled dynamics.

The steady-state model is based on the linear bicycle model, which relates the wheel angle $\delta$ to the yaw rate. The complete description is shown in Equation 17, where $\dot{\psi}_{ss}$ is the steady-state yaw rate, $u$ is the vehicle's longitudinal velocity, $L$ is the wheelbase, $K_{us}$ is the understeer gradient, and $g$ is the gravitational acceleration.

$$\dot{\psi}_{ss} = \frac{u}{L + \frac{K_{us}u^2}{g}}\delta \tag{17}$$

The linear bicycle model's predictive capability is limited since it does not account for the nonlinearities of tire grip. Tires generate a finite force, therefore this limitation must be considered in the reference generator. The lateral acceleration $a_y$ must remain within the grip limit, so $a_y \leq \mu g$, where $\mu$ is the friction coefficient. Because the controller manages yaw rate, $a_{y_{\max}}$ must be converted into $\dot{\psi}_{\max}$. Lateral acceleration $a_y$ can be expressed as the product of longitudinal velocity and yaw rate, combined with the second derivative of the y-position. To account for potential longitudinal acceleration, it is assumed that approximately 85% of the total grip is available for lateral force generation. Consequently, the bounded reference yaw rate is described by Equation 18.

$$\begin{aligned} \dot{\psi}_{max} &= 0.85 \cdot \frac{\mu g}{u} \\ \dot{\psi}_{sat} &= \begin{cases} \dot{\psi}_{ss}, & |\dot{\psi}_{ss}| \leq |\dot{\psi}_{max}| \\ \pm\dot{\psi}_{max}, & |\dot{\psi}_{ss}| > |\dot{\psi}_{max}| \end{cases} \end{aligned} \tag{18}$$

The final step involves applying a second-order transfer function to the bounded yaw signal to capture unmodeled dynamics, such as weight transfer. The transfer function used is presented in Equation 19. For the test vehicle, the natural frequency $\omega_0 = 11$ rad/s, the damping ratio was set to $\zeta = 0.7$, and the time constant $\tau = 0.09$ s.

$$\dot{\psi}_{ref} = \frac{\omega_0^2\,(1+\tau s)}{s^2 + 2\zeta\omega_0 s + \omega_0^2}\dot{\psi}_{sat} \tag{19}$$

To utilize the LQR controller, a linear state-space formulation of the relationship between the steering angle and yaw rate needs to be established. The bicycle model was employed again, assuming a constant longitudinal velocity. Consequently, the $A$ and $B$ matrices are derived, as shown in Equation 20, with $l_f$ and $l_r$ referring to the distance between the center of gravity and the front and rear axles, $C_{\alpha f}$ and $C_{\alpha r}$ standing for front and rear cornering stiffness, and $I_z$ representing the mass moment of inertia around the vertical axis. The formulation is velocity-dependent, so the optimal gain $K$ is calculated for several discrete velocities and then linearly interpolated to create a gain-scheduled proportional controller with optimal gains. The optimal gain $K$ can be calculated by solving the Riccati equation, using the `lqr` function in MATLAB.

$$A(u) = \left[\frac{l_r^2 C_{\alpha r} + l_f^2 C_{\alpha f}}{I_z V_{ref}}\right]$$
$$B(u) = \left[\frac{l_f C_{\alpha f}}{I_z}\right] \tag{20}$$

**E. Testing Environment**

The CBVF safety controller will be evaluated on four different road conditions and with two different base controllers, using a modified Sine and Dwell test. The Sine and Dwell test is an open-loop steering command test that is the standard for stability testing. The steering maneuver consists of three phases. First, three-fourths of a period of a sine wave (initial left-to-right movement). Then, the steering is held at the same position for a given amount of time (the dwell). Finally, the steering is recovered to the neutral position, following a quarter-period sine wave (steer centering). The exact steering deflection used is shown in Figure 2.

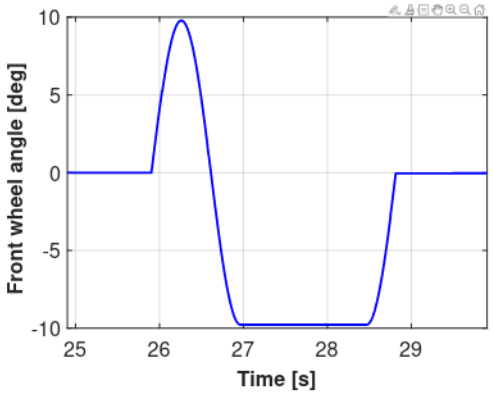


**Fig. 2 Example time trace of the front steering angle**

The test will be conducted on a high grip surface ($\mu = 1.0$) at 100 km/h, a low grip surface ($\mu = 0.2$) at 50 km/h, and two transitions from high to low grip conditions at 70 km/h. The "late" $\mu$ transition happens at the beginning of the dwell portion of the steering input. The "early" $\mu$ transition happens in the middle of the initial sine input. The procedures are summarized in Table 2.

| Test Scenario | Initial $\mu$ | Final $\mu$ | Speed [km/h] | Transition Timing |
|---|---|---|---|---|
| High Grip | 1.0 | - | 100 | - |
| Low Grip | 0.2 | - | 50 | - |
| Early Transition | 1.0 | 0.2 | 70 | Middle of initial sine |
| Late Transition | 1.0 | 0.2 | 70 | Beginning of dwell |

**Table 2 Testing scenario specifications for modified Sine and Dwell test**

These scenarios are not the typical part of ESC testing. Per regulations, only high-grip testing needs to be performed. However, $\mu$ transitions are a challenging case where the differences in stability controllers should be evident. The testing scenarios are summarized in.

As mentioned in subsection III.D, the safety controller will be tested in combination with two different controllers. One very important point is that all controllers assume high grip conditions. This means that in the computation of the reference yaw rates for the LQR controller and the BRT computation in the CBVF controller, $\mu$ is assumed to be 1.0. This is done since online estimation of the available grip level is difficult and remains an open research question [32]. Additionally, if the controllers can perform their stabilizing duties without the need to be tuned for different driving scenarios, it would vastly simplify the creation of stability control systems.

## IV. Results

The results are split into two parts. First, the parameters for the BRT computation are evaluated. This is where the effects of time horizon, relaxation parameter, and maximum applied moment are studied. The goal of this study is to understand what effects these parameters have on the structure of the BRT and the performance of the safety controller. The second part focuses on the performance comparison between the baseline controller and the safety controller.

### A. Controller Tuning

This parameter study investigates how BRT design parameters affect safety controller performance when implemented on a realistic vehicle model. Three maximum moment values were selected: $M_{z_{\text{MAX}}} = 1$ kNm, 5 kNm, and 10 kNm, representing different theoretical control authorities. Three time horizons were also selected: $t_{\text{end}} = 0.2$ s, 0.4 s, and 0.6 s, representing different prediction horizons. For each parameter combination, the relaxation parameter $\gamma$ was tuned to achieve optimal performance in both high- and low-grip conditions.

The $M_{z_{\text{MAX}}}$ values serve as design parameters for exploring different safe set geometries in the BRT computation rather than representations of physical actuator limits. Since the simplified bicycle model uses a pure moment input while the CarMaker validation employs differential braking, no direct correspondence exists between the theoretical $M_{z_{\text{MAX}}}$ and the actual system control authority. This parameter study, therefore, investigates how different theoretical safe sets perform when implemented on a realistic vehicle model.

The time horizon $t_{\text{end}}$ is a design parameter that must be chosen for the BRT computation. It defines the recovery time window within which the system can return to the target set. For example, with $t_{\text{end}} = 0.2$ s, states inside the BRT can reach the target set within 0.2 s. Since no established time windows exist for vehicle stability control, this parameter must be determined empirically.

Each combination of $M_{z_{\text{MAX}}}$ and $t_{\text{end}}$ requires individual tuning of $\gamma$, as the value functions have different magnitudes for each configuration. To comprehensively evaluate the influence of $\gamma$ on each safety controller, simulations were performed on both high- and low-grip surfaces. The optimal $\gamma$ for each configuration was identified through a two-stage process: first, a coarse grid search tested 35 equally spaced values between 1 and 200; then, a Newton-Raphson method refined the tuning with a convergence tolerance of 0.01.

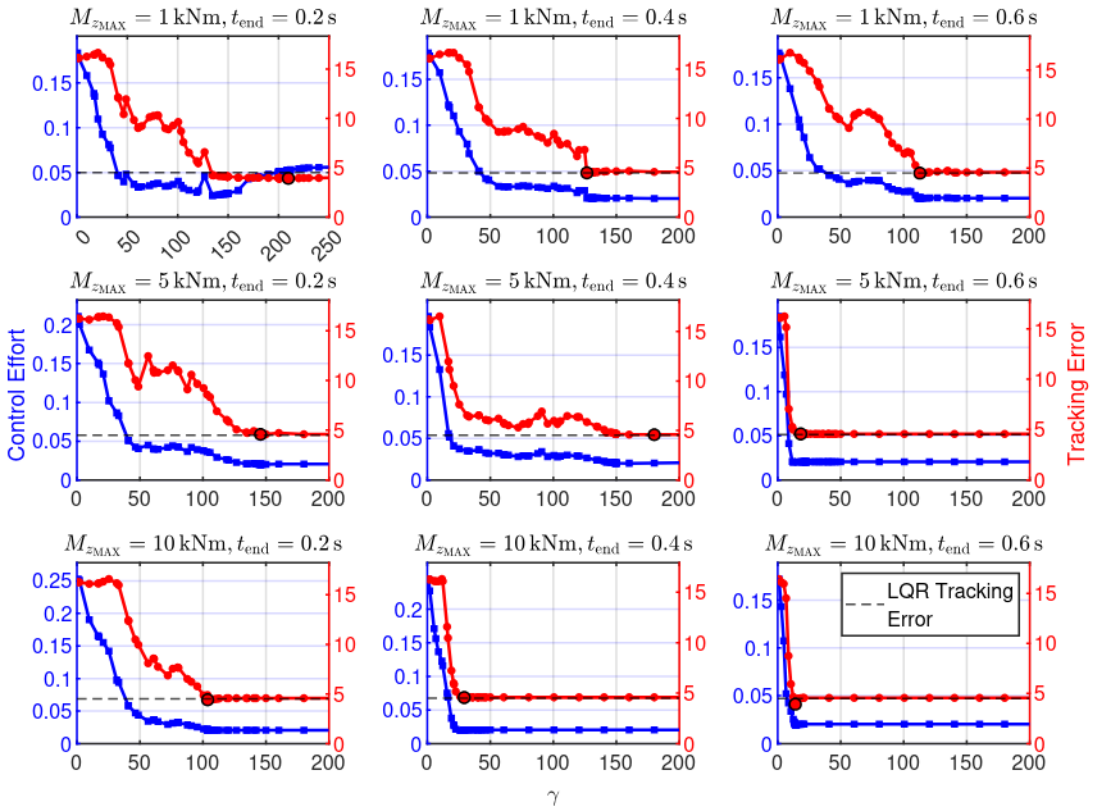


**Fig. 3 Results of $\gamma$ sweep on a high grip surface**

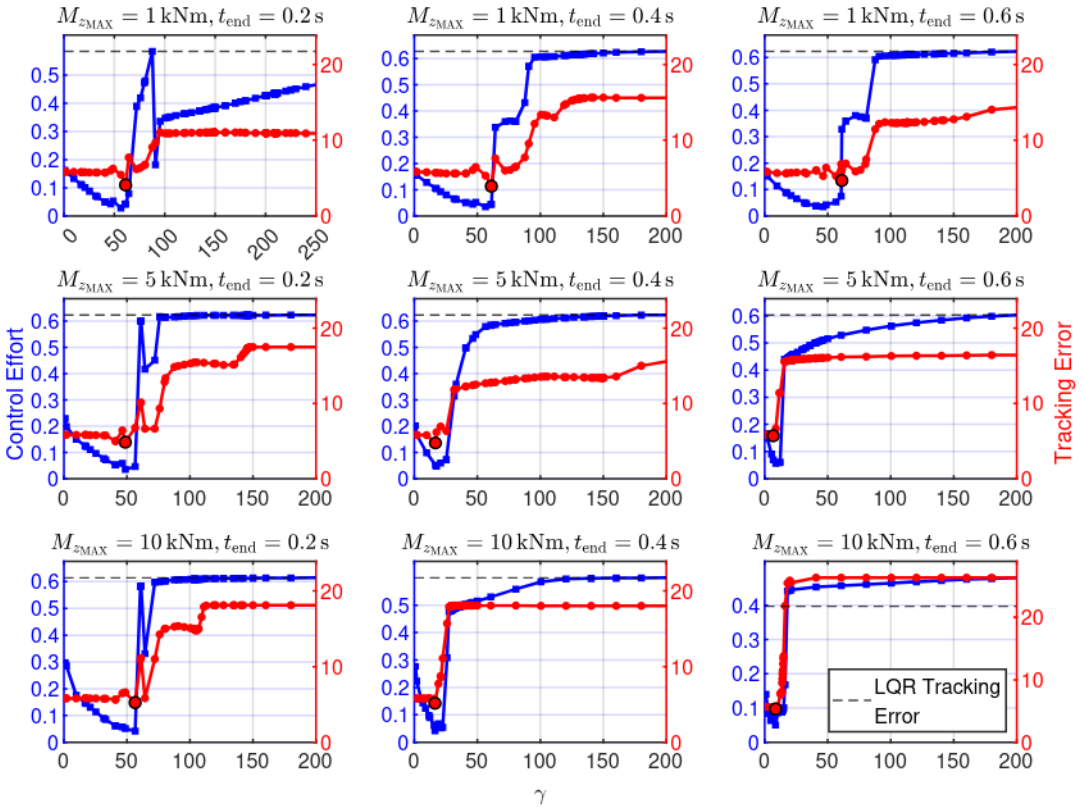


**Fig. 4 Results of $\gamma$ sweep on a low grip surface**

Tracking performance is evaluated using the root mean square error (RMSE) between the achieved simulation yaw rate and the reference yaw rate. For high-grip conditions, the reference is the LQR-computed trajectory. For low-grip conditions, the reference is the grip-available line, computed using the same method as the LQR reference but with dynamically measured $\mu$ values. Control effort is quantified as the RMSE between the applied control moment and zero.

In the QP solver implementation, control bounds were set to $10^5$ Nm to avoid artificial constraints unrelated to the CBVF methodology. This approach isolates the safety filter's contribution from arbitrary actuator management effects, enabling direct assessment of how different BRT parameterizations affect controller performance.

An important aspect of this evaluation is that both the baseline LQR controller and all safety filter configurations were used, assuming high-grip conditions. The LQR reference trajectory assumes maximum available friction, and the safety filters were computed assuming a high $\mu$. This experimental design specifically evaluates the controllers' ability to maintain acceptable performance across varying surface conditions without requiring recalibration. This addresses a significant practical challenge in automotive control systems, where electronic control units traditionally require extensive tuning across numerous surface types, environmental conditions, and vehicle loading scenarios. A controller that performs well under its design conditions while gracefully degrading or maintaining safety on other surfaces represents a substantial improvement over conventional approaches that necessitate separate tuning for each operating regime.

All controller configurations were validated against ECE R140 sine dwell test requirements on dry asphalt [33]. The yaw rate recovery criteria require that the yaw rate must not exceed 35% of the peak value at 1 second and 20% of the peak value at 1.75 seconds after completion of the steering input. The lateral displacement requirement mandates at least 1.83 m measured at 1.07 seconds after the Beginning of Steer. The only safety filters that failed to meet these criteria were those with very low $\gamma$ tunings.

All configurations exhibit an initial plateau at low $\gamma$ values, with tracking errors exceeding 15. These safety controllers are excessively conservative, restricting lateral motion to the extent that they fail to meet the lateral displacement criteria. However, all tunings near the optimal $\gamma$ successfully pass the regulatory requirements.

The results of the $\gamma$ sweeps are presented in Figure 6 and Figure 7, with the best-performing $\gamma$ for each configuration compared to the baseline controller in Table 3 and Table 4.

| Method | Tracking Error | vs Baseline | Control Effort | vs Baseline | Gamma |
|---|---|---|---|---|---|
| LQR w/ SF - Baseline | 4.5063 | – | 0.0202 | – | – |
| $M_{z_{MAX}}$=1kNm, $t_{end}$=0.2s | 3.9452 | -12.45% | 0.0415 | +106.12% | 209.74 |
| $M_{z_{MAX}}$=1kNm, $t_{end}$=0.4s | 4.5044 | -0.04% | 0.0166 | -17.88% | 126.63 |
| $M_{z_{MAX}}$=1kNm, $t_{end}$=0.6s | 4.4996 | -0.15% | 0.0164 | -18.72% | 112.84 |
| $M_{z_{MAX}}$=5kNm, $t_{end}$=0.2s | 4.5778 | +1.59% | 0.0158 | -21.45% | 145.74 |
| $M_{z_{MAX}}$=5kNm, $t_{end}$=0.4s | 4.5653 | +1.31% | 0.0161 | -20.16% | 180.10 |
| $M_{z_{MAX}}$=5kNm, $t_{end}$=0.6s | 4.5799 | +1.63% | 0.0162 | -19.65% | 18.33 |
| $M_{z_{MAX}}$=10kNm, $t_{end}$=0.2s | 4.4172 | -1.98% | 0.0166 | -17.86% | 103.72 |
| $M_{z_{MAX}}$=10kNm, $t_{end}$=0.4s | 4.5759 | +1.55% | 0.0159 | -21.19% | 29.30 |
| $M_{z_{MAX}}$=10kNm, $t_{end}$=0.6s | 3.9423 | -12.51% | 0.0160 | -20.77% | 14.00 |

**Table 3 Control performance results for optimal $\gamma$ (dry conditions)**

| Method | Tracking Error | vs Baseline | Control Effort | vs Baseline | Gamma |
|---|---|---|---|---|---|
| LQR w/ SF - Baseline | 21.7343 | – | 0.1242 | – | – |
| $M_{z_{MAX}}$=1kNm, $t_{end}$=0.2s | 4.0815 | -81.22% | 0.0339 | -72.69% | 61.20 |
| $M_{z_{MAX}}$=1kNm, $t_{end}$=0.4s | 3.9063 | -82.03% | 0.0357 | -71.26% | 61.20 |
| $M_{z_{MAX}}$=1kNm, $t_{end}$=0.6s | 4.6759 | -78.49% | 0.2596 | +109.05% | 61.20 |
| $M_{z_{MAX}}$=5kNm, $t_{end}$=0.2s | 4.8414 | -77.72% | 0.0275 | -77.90% | 48.74 |
| $M_{z_{MAX}}$=5kNm, $t_{end}$=0.4s | 4.7391 | -78.20% | 0.0404 | -67.46% | 16.80 |
| $M_{z_{MAX}}$=5kNm, $t_{end}$=0.6s | 5.7104 | -73.73% | 0.0551 | -55.62% | 6.80 |
| $M_{z_{MAX}}$=10kNm, $t_{end}$=0.2s | 5.2309 | -75.93% | 0.0328 | -73.57% | 56.53 |
| $M_{z_{MAX}}$=10kNm, $t_{end}$=0.4s | 5.1520 | -76.30% | 0.0333 | -73.19% | 16.40 |
| $M_{z_{MAX}}$=10kNm, $t_{end}$=0.6s | 5.3506 | -75.38% | 0.0403 | -67.59% | 8.60 |

**Table 4 Control performance results for optimal $\gamma$ (wet conditions)**

The parameter sweep reveals several important characteristics of how BRT design parameters affect safety controller behavior. Most controller configurations eventually converge to pass-through behavior, where performance becomes identical to the baseline LQR controller with no safety intervention occurring. This is evident from the tracking error and control effort metrics matching the baseline values. Only two configurations demonstrate improved tracking performance on high-grip surfaces: the $M_{z_{MAX}}$ = 1 kNm with $t_{end}$ = 0.2 s configuration and the $M_{z_{MAX}}$ = 10 kNm with $t_{end}$ = 0.6 s configuration, both achieving approximately 12.5% reduction in tracking error compared to the baseline.

A notable distinction between these two configurations lies in their sensitivity to the relaxation parameter $\gamma$. As shown in Figure 6, the $M_{z_{MAX}}$ = 1 kNm, $t_{end}$ = 0.2 s configuration exhibits a broad region of near-optimal performance, with the tracking error remaining low across a wide range of $\gamma$ values (approximately 150-250). In contrast, the $M_{z_{MAX}}$ = 10 kNm, $t_{end}$ = 0.6 s configuration displays an extremely narrow performance peak, with optimal behavior occurring only within a small window around $\gamma$ = 14.

This difference in $\gamma$ sensitivity suggests fundamentally different controller characteristics. The wide optimal region of the $M_{z_{MAX}}$ = 1 kNm, $t_{end}$ = 0.2 s configuration resembles traditional safety controllers that maintain consistent protective behavior across a range of tuning parameters. Conversely, the sharp peak exhibited by the $M_{z_{MAX}}$ = 10 kNm, $t_{end}$ = 0.6 s configuration indicates operation as a Hamilton-Jacobi-informed quasi-steady-state controller, where small changes in $\gamma$ produce significant behavioral shifts. This fundamental difference in controller behavior is illustrated in Figure 5, which shows the safe set geometries and vehicle trajectories in state space. The $M_{z_{MAX}}$ = 1 kNm configuration (Figure 5a) exhibits a more constrained safe set, resulting in tighter trajectory regulation and more conservative intervention. In contrast, the $M_{z_{MAX}}$ = 10 kNm configuration (Figure 5b) provides a larger safe set volume, allowing greater state excursions while maintaining safety constraints.

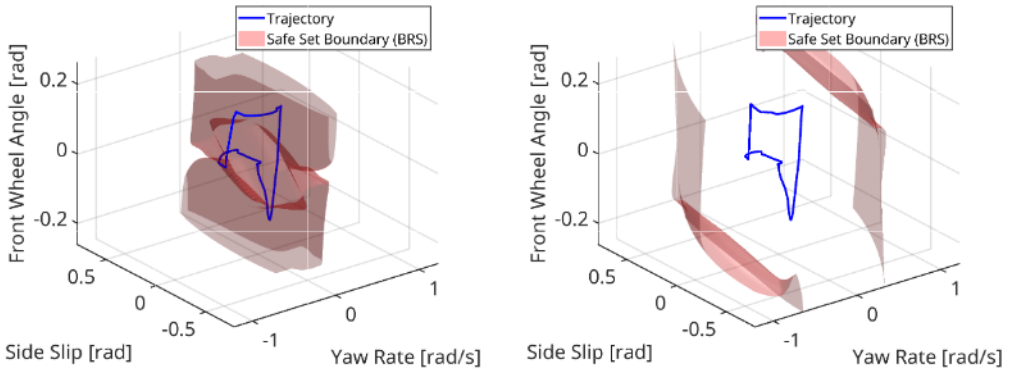


**(a) Safety controller based on $M_{z_{\mathrm{MAX}}}$ = 1kNm, $t_{\mathrm{end}}$ = 0.2s with $\gamma$ = 150 on low grip surface**

**(b) Safety controller based on $M_{z_{\mathrm{MAX}}}$ = 10kNm, $t_{\mathrm{end}}$ = 0.6s with $\gamma$ = 13.1 on low grip surface**

**Fig. 5 Example BRTs with trajectories of tests performed on a low grip surface**

The optimal $\gamma$ values for high-grip and low-grip conditions diverge significantly for configurations with shorter time horizons or smaller maximum moments, as demonstrated in Figure 3 and Figure 4. The $M_{z_{\mathrm{MAX}}}$ = 10 kNm, $t_{\mathrm{end}}$ = 0.6 s configuration represents a notable exception, achieving optimal performance in both grip conditions with similar $\gamma$ values (14.00 for dry, 8.60 for wet). However, this convergence may be attributable to the configuration's high $\gamma$ sensitivity, where substantial behavioral changes occur with relatively small parameter adjustments.

On low-grip surfaces, nearly all controller configurations substantially outperform the baseline LQR controller, achieving 70-82% reductions in tracking error as detailed in Table 4. This improvement primarily stems from the safety filters' ability to prevent vehicle spinout by constraining the system within the safe set. The baseline LQR controller, lacking this constraint, attempts to track the reference trajectory beyond the vehicle's physical capabilities, resulting in instability and large tracking errors.

The time-domain behavior illustrated in Figure 6 and Figure 7 reveals important characteristics of the safety filter interventions. On the high-grip surface, the safety-filtered controllers apply only minimal corrective moments compared to the baseline LQR. These subtle interventions—barely visible in the moment traces—nonetheless produce substantial improvements in tracking performance, with up to 12.5% reduction in yaw rate error. This demonstrates the safety filters' ability to make targeted, efficient corrections that meaningfully improve vehicle behavior without excessive control authority.

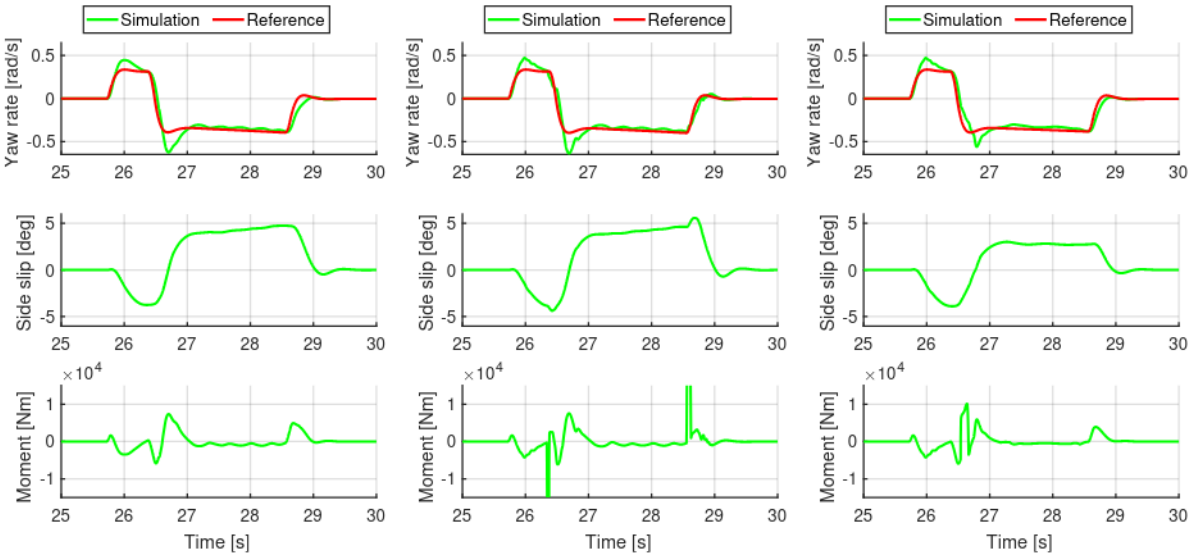


**(a) LQR with no safety controller**

**(b) LQR with safety controller computed on $M_{z_{\mathrm{MAX}}}$=1kNm and $t_{\mathrm{end}}$=0.2s with $\gamma$=209.74**

**(c) LQR with safety controller computed on $M_{z_{\mathrm{MAX}}}$=10kNm and $t_{\mathrm{end}}$=0.6s with $\gamma$=14.00**

**Fig. 6 Controller performance evaluation on a high grip surface**

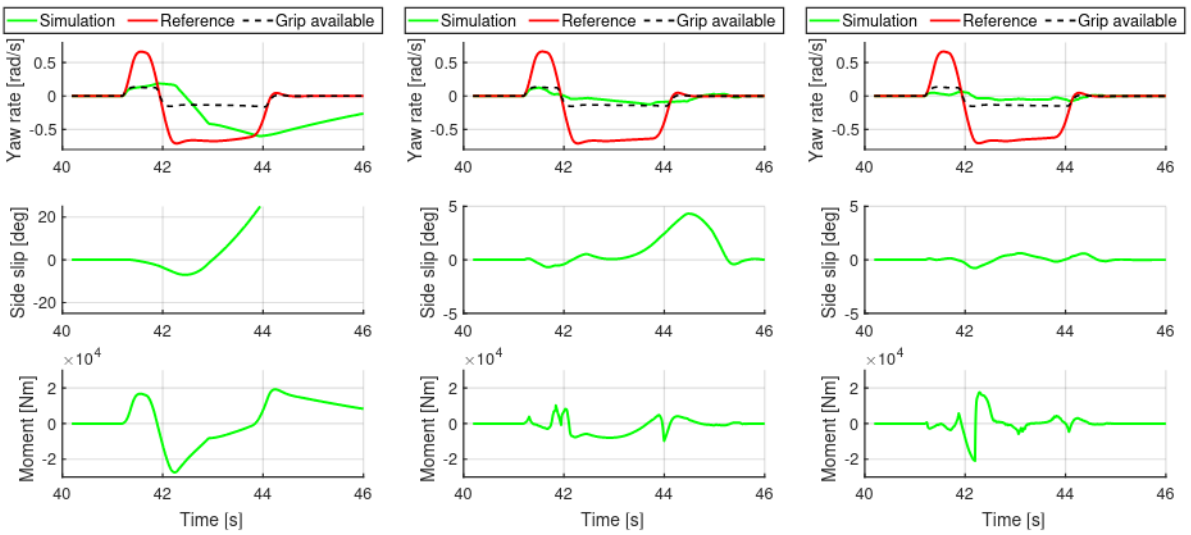


**(a) LQR with no safety controller**

**(b) LQR with safety controller computed on $M_{z_{MAX}}$=1kNm and $t_{end}$=0.2s with $\gamma$=61.20**

**(c) LQR with safety controller computed on $M_{z_{MAX}}$=10kNm and $t_{end}$=0.6s with $\gamma = 8.60$**

**Fig. 7 Controller performance evaluation on a low grip surface**

In contrast, the low-grip scenarios show dramatically different behavior. Figure Figure 7(a) clearly shows the baseline LQR controller causing vehicle spinout, with the yaw rate substantially exceeding the grip-available limit and large sideslip angles developing. The safety-filtered controllers successfully prevent this instability by constraining the vehicle states within the safe set. However, an important limitation emerges: the $\gamma$ values optimized for high-grip conditions prove overly conservative when applied to low-grip surfaces. This is particularly evident in the reduced sideslip excursions and more cautious yaw rate profiles compared to what the available grip could theoretically support. The controllers prioritize safety over performance, maintaining larger margins from the stability boundaries than strictly necessary.

This grip-dependent behavior highlights a fundamental challenge in implementing these controllers: optimal $\gamma$ values differ significantly between surface conditions. While the $M_{z_{MAX}} = 10$ kNm, $t_{end} = 0.6$ s configuration shows the closest alignment between optimal $\gamma$ values across conditions (14.00 for dry, 8.60 for wet), even this configuration requires different tuning for optimal performance. In practice, a compromise $\gamma$ value must be selected that balances performance across the expected operating envelope. This trade-off represents an inherent limitation of the one-shot tuning approach—while a single controller can operate safely across conditions, achieving optimal performance in all scenarios simultaneously remains elusive. The choice of compromise $\gamma$ depends on the desired balance between high-grip performance optimization and low-grip conservatism.

Based on the parameter sweep results, two controller configurations with compromise $\gamma$ values were selected for final validation. For the $M_{z_{MAX}} = 10$ kNm, $t_{end} = 0.6$ s configuration, $\gamma = 13.1$ was chosen as the compromise value. This selection corresponds to the lowest $\gamma$ within the narrow high-grip performance peak while avoiding the rapidly degrading performance region in low-grip conditions. As shown in Figure 7, $\gamma$ values below approximately 13 result in sharply increasing tracking errors on wet surfaces, making 13.1 a pragmatic lower bound that maintains acceptable performance across both conditions.

The selection process for the $M_{z_{MAX}} = 1$ kNm, $t_{end} = 0.2$ s configuration proved more challenging due to the lack of overlap between optimal performance regions. A compromise value of $\gamma = 150$ was selected, positioned at the lower end of the broad high-grip performance plateau visible in Figure 6. This choice prioritizes two considerations: it maintains near-optimal tracking performance on dry surfaces while avoiding the excessive control effort observed at higher $\gamma$ values in wet conditions, as evidenced in Figure 4. The resulting performance of both configurations under their compromise $\gamma$ tunings is illustrated in Figure 8, demonstrating that both controllers successfully prevent instability across grip conditions while maintaining reasonable tracking performance where the vehicle dynamics permit.

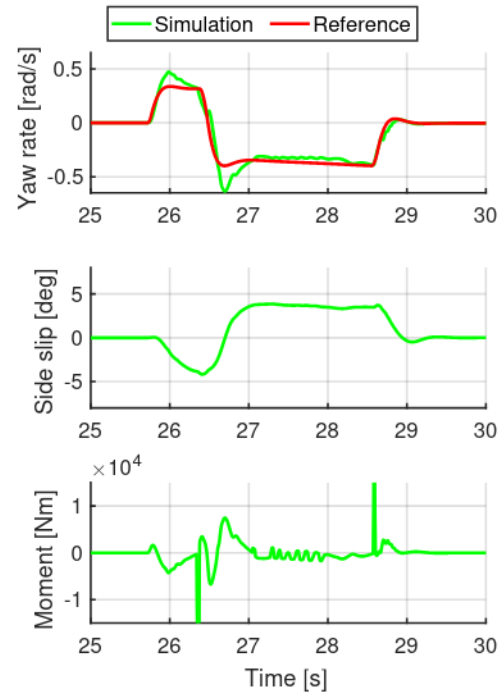


**(a) Safety controller based on $M_{z_{MAX}}$ = 1kNm, $t_{end}$ = 0.2s with $\gamma$ = 150 on high grip surface**

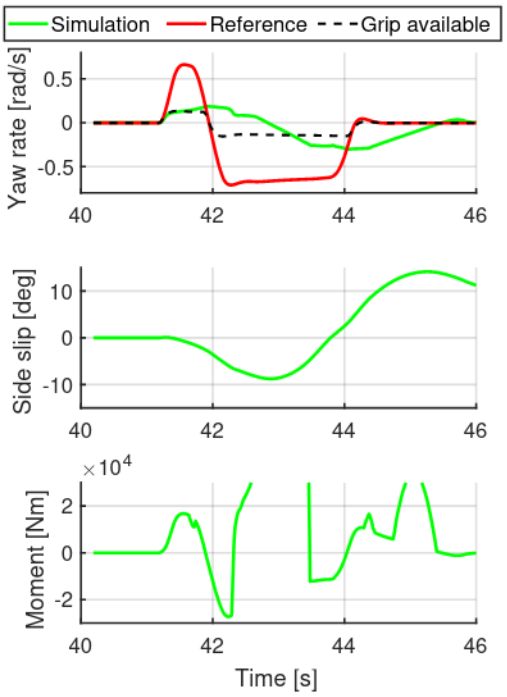


**(b) Safety controller based on $M_{z_{MAX}}$ = 1kNm, $t_{end}$ = 0.2s with $\gamma$ = 150 on low grip surface**

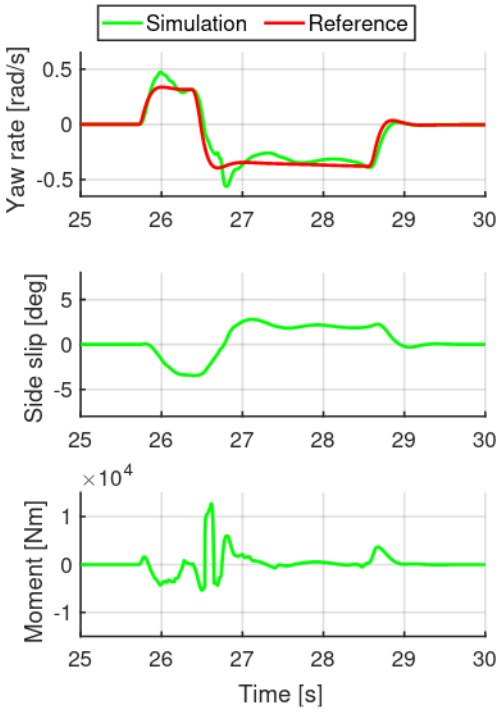


**(c) Safety controller based on $M_{z_{MAX}}$ = 10kNm, $t_{end}$ = 0.6s with $\gamma$ = 13.1 on high grip surface**

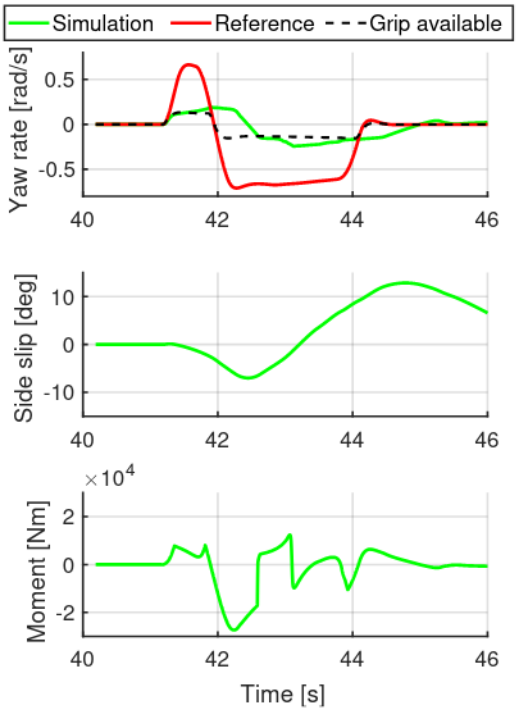


**(d) Safety controller based on $M_{z_{MAX}}$ = 10kNm, $t_{end}$ = 0.6s with $\gamma$ = 13.1 on low grip surface**

**Fig. 8 Safety controller performance evaluation with final tuning of $\gamma$**

**B. Controller Comparison**

The final validation compares the two selected controller configurations, $M_{z_{\text{MAX}}} = 1$ kNm with $t_{\text{end}} = 0.2$ s at $\gamma = 150$ and $M_{z_{\text{MAX}}} = 10$ kNm with $t_{\text{end}} = 0.6$ s at $\gamma = 13.1$, against the baseline LQR controller across four surface conditions and three steering amplitudes. The conditions include high-grip dry asphalt at 100 km/h, low-grip wet asphalt at 50 km/h, and two grip transition scenarios at 70 km/h where the surface changes from high to low grip either during initial turning (early transition) or during the dwell phase (late transition). The results are summarized in Table 5 and Table 6, with representative time traces shown in Figure 9 and Figure 10.

On high-grip surfaces, both controllers demonstrate minimal intervention behavior consistent with their design as safety filters rather than performance optimizers. The $M_{z_{\text{MAX}}} = 10$ kNm configuration exhibits largely pass-through behavior, with tracking performance within 10% of the baseline and only modest increases in control effort. The $M_{z_{\text{MAX}}} = 1$ kNm configuration shows slightly more active intervention, achieving 5-7% tracking improvements at lower steering amplitudes but degrading performance by 14% at the highest amplitude. As illustrated in Figure 9a and Figure 10a, the applied moments remain small and the safety filter allows the baseline controller to operate largely unimpeded, making only subtle corrections to maintain the vehicle within the safe set.

The low-grip scenarios reveal fundamentally different characteristics between the two controllers. The $M_{z_{\text{MAX}}} = 10$ kNm configuration demonstrates robust performance, achieving 68-71% tracking error reduction at higher steering amplitudes while maintaining reasonable control effort. In contrast, the $M_{z_{\text{MAX}}} = 1$ kNm configuration achieves similar tracking improvements (53-58%) but at the cost of substantially increased control effort, particularly at the 170° amplitude, where control effort increases by 174%. The time traces in Figure 9b show both configurations successfully preventing the spinout that occurs with the baseline LQR controller. However, the moment plots reveal that the $M_{z_{\text{MAX}}} = 1$ kNm controller applies larger corrective moments more frequently, consistent with its tighter safe set constraints observed in Figure 5a.

The grip transition scenarios expose a critical limitation of the $M_{z_{\text{MAX}}} = 1$ kNm, $t_{\text{end}} = 0.2$ s configuration. During early transitions, where the surface changes from high to low grip during initial turning, this controller exhibits catastrophic performance degradation, with tracking errors increasing by 1.9-2.3x compared to the baseline and control effort surging by 5.3-6.6x. The extreme control moments visible in Figure 9c indicate that the vehicle exits the BRT during the transition, triggering the safety filter to request the maximum allowed control moment in the QP solver. The narrow safe set geometry proves insufficient to accommodate the rapid dynamic changes during grip transitions, resulting in vehicle spinout despite the safety filter's intervention.

In contrast, the $M_{z_{\text{MAX}}} = 10$ kNm, $t_{\text{end}} = 0.6$ s configuration maintains stable operation throughout all transition scenarios. For early transitions, it achieves 9-70% tracking improvements with modest control effort increases of 5-49%. The late transition results are similarly favorable, with 35-80% tracking improvements. The longer time horizon and larger safe set volume provide sufficient margin to accommodate transient dynamics during grip changes, as evidenced by the smooth moment traces in Figure 10c and Figure 10d. Notably, no spinouts occur with this configuration under any tested condition.

| Surface Condition | Steering Amplitude | LQR | | LQR W/ SF | | Improvement | |
|---|---|---|---|---|---|---|---|
| | | Tracking Error | Control Effort | Tracking Error | Control Effort | Tracking | Control |
| **High-grip** | 100° | 4.11 | 0.015 | 3.88 | 0.020 | 5.6% | -33.3% |
| | 170° | 4.51 | 0.016 | 4.20 | 0.020 | 6.7% | -20.8% |
| | 250° | 4.97 | 0.018 | 5.66 | 0.018 | -14.1% | -0.3% |
| **Low-grip** | 100° | 10.66 | 0.074 | 12.52 | 0.061 | -17.5% | 17.4% |
| | 170° | 26.31 | 0.111 | 11.01 | 0.305 | 58.1% | -174.2% |
| | 250° | 21.56 | 0.109 | 10.03 | 0.201 | 53.5% | -84.4% |
| **Early Transition** | 100° | 13.85 | 0.105 | 39.46 | 0.808 | -185.0% | -666.2% |
| | 170° | 15.76 | 0.120 | 51.36 | 0.754 | -225.8% | -530.3% |
| | 250° | 14.82 | 0.117 | 47.81 | 0.783 | -222.5% | -570.5% |
| **Late Transition** | 100° | 22.34 | 0.059 | 12.74 | 0.710 | 42.9% | -1108.2% |
| | 170° | 37.78 | 0.118 | 15.55 | 0.786 | 58.8% | -565.9% |
| | 250° | 40.06 | 0.126 | 13.89 | 0.750 | 65.3% | -494.3% |

**Table 5 Controller performance comparison across surface conditions ($M_z$=1kNm, t=-0.2s, $\gamma$=150)**

| Surface Condition | Steering Amplitude | LQR | | LQR W/ SF | | Improvement | |
|---|---|---|---|---|---|---|---|
| | | Tracking Error | Control Effort | Tracking Error | Control Effort | Tracking | Control |
| **High-grip** | 100° | 4.11 | 0.015 | 4.52 | 0.020 | -10.0% | -33.4% |
| | 170° | 4.51 | 0.016 | 4.22 | 0.019 | 6.4% | -14.5% |
| | 250° | 4.97 | 0.018 | 6.01 | 0.020 | -21.1% | -13.6% |
| **Low-grip** | 100° | 10.66 | 0.074 | 10.66 | 0.073 | 0.0% | 1.2% |
| | 170° | 26.31 | 0.111 | 7.71 | 0.070 | 70.7% | 37.5% |
| | 250° | 21.56 | 0.109 | 6.82 | 0.080 | 68.4% | 26.3% |
| **Early Transition** | 100° | 13.85 | 0.105 | 12.60 | 0.101 | 9.0% | 4.6% |
| | 170° | 15.76 | 0.120 | 8.31 | 0.062 | 47.3% | 48.6% |
| | 250° | 14.82 | 0.117 | 4.47 | 0.067 | 69.8% | 42.6% |
| **Late Transition** | 100° | 22.34 | 0.059 | 14.60 | 0.154 | 34.6% | -162.5% |
| | 170° | 37.78 | 0.118 | 11.89 | 0.048 | 68.5% | 59.2% |
| | 250° | 40.06 | 0.126 | 8.14 | 0.061 | 79.7% | 51.8% |

**Table 6 Controller performance comparison across surface conditions ($M_z$=10kNm, t=-0.6s, $\gamma$=13.1)**

An important observation across select scenarios is the occurrence of extreme control moment requests when the vehicle state exits the BRT. These moments, causing the QP solver to hit the 1e5 Nm limits, represent the safety filter's attempt to recover stability from states outside the computed safe set. While the $M_{z_{MAX}} = 10$ kNm configuration rarely encounters this situation, the $M_{z_{MAX}} = 1$ kNm configuration frequently exits the BRT during transitions, leading to aggressive and often unsuccessful recovery attempts.

This behavior highlights a potentially interesting property of the CBVF method. In principle, since $\mu$ is part of the system dynamics used in the HJ calculation, it is expected that it wouldn't perform well on surfaces with a different $\mu$. Therefore, even though the $M_{z_{MAX}} = 10$ kNm set up is likely less representative of the real safety bounds of a passenger vehicle, its broader BRT includes gradient data, which increases the performance of the safety controller. This configuration and tuning turned the CBVF controller from a boundary-obeying filter to an HJ gradient-informed augmentation system. The main objective is not to stay within the boundary of the safe set, but to stop the system from diverging too quickly from the steady-state behavior. The mathematical formulation remains the same, but the choice of the parameters for the HJ solve and the tuning of $\gamma$ change the nature of the controller. This highlights an unexplored and potentially promising direction to develop HJ-based controllers towards. The value function contains a wealth of information about the system. Most methods have focused on the zero-level set of the value function, and in the context of safety filters, it is irreplaceable, but there is clearly more information to be gathered from HJ approaches, and this experiment, on order-reduced systems, demonstrates that clearly.

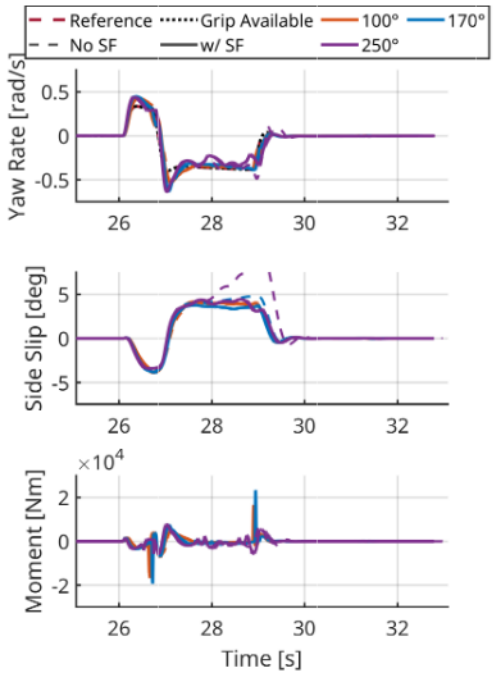


(a) High grip surface at 100 km/h

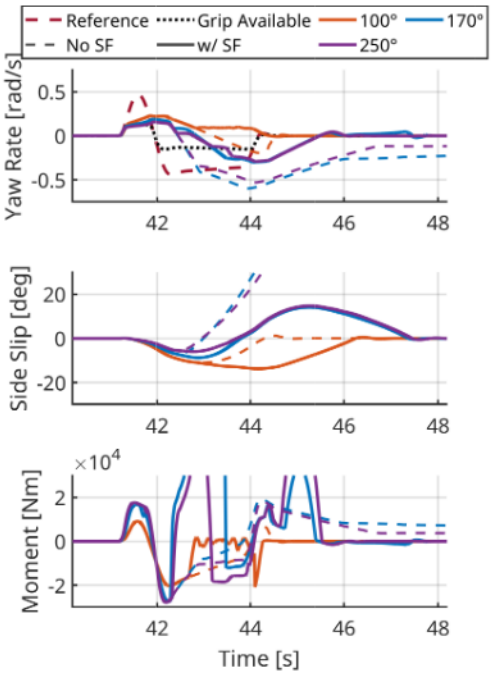


(b) Low grip surface at 50 km/h

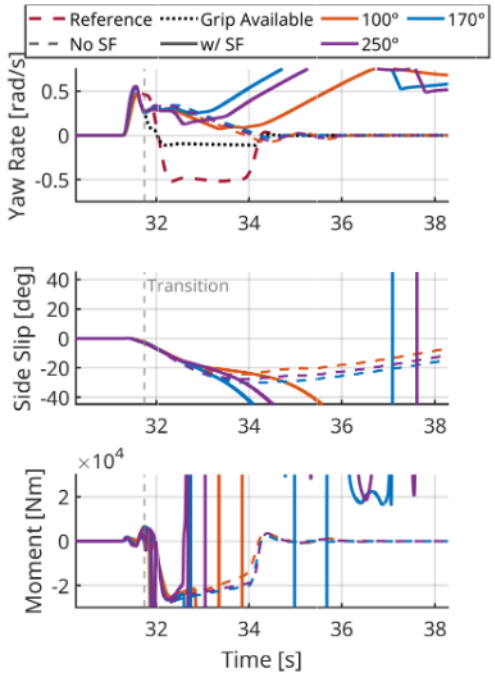


(c) High to low grip transition during initial turning at 70 km/h

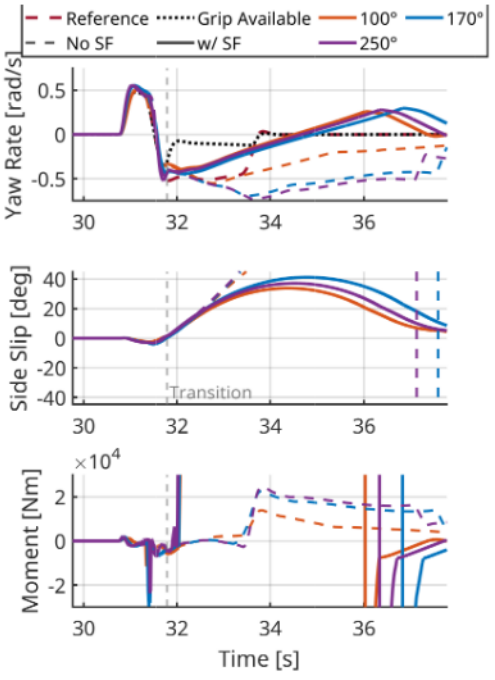


(d) High to low grip transition during dwell at 70 km/h

**Fig. 9 Comparison of LQR with the CBVF safety controller enabled ($M_z$=1kNm, t=-0.2s, $\gamma$=150)**

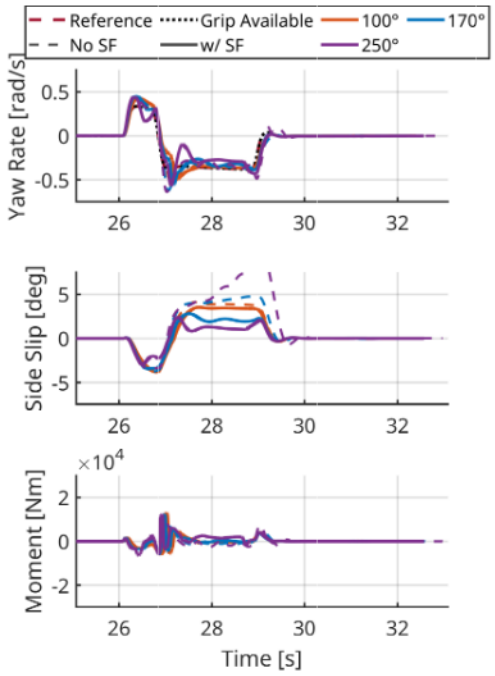


(a) High grip surface at 100 km/h

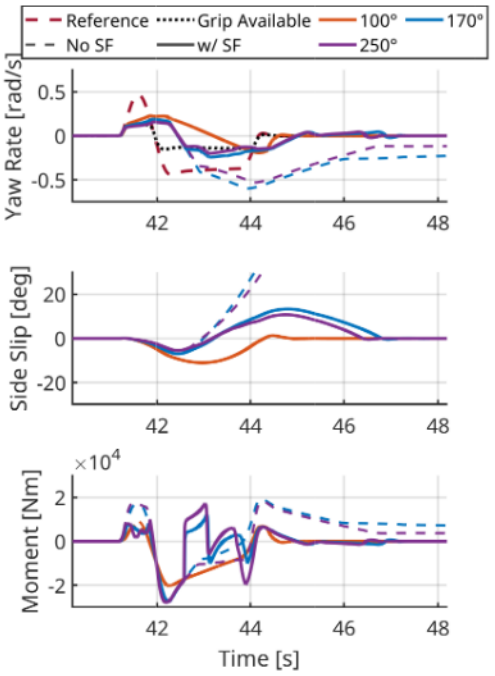


(b) Low grip surface at 50 km/h

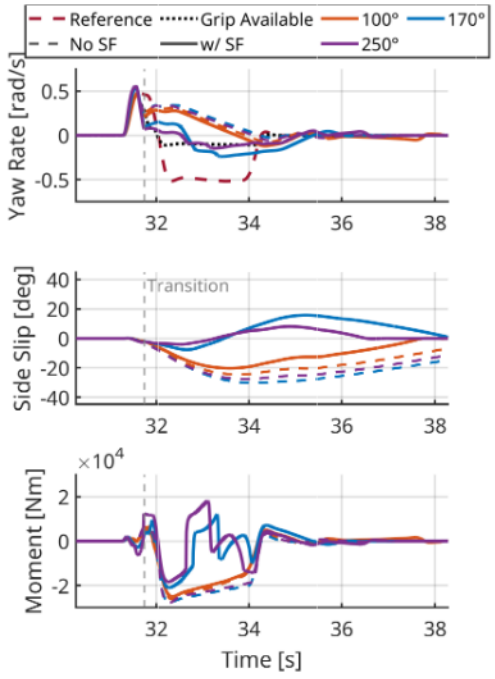


(c) High to low grip transition during initial turning at 70 km/h

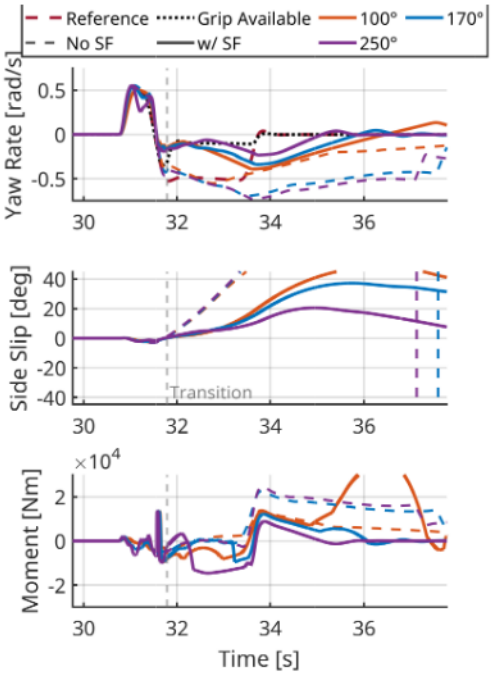


(d) High to low grip transition during dwell at 70 km/h

**Fig. 10 Comparison of LQR with the CBVF safety controller enabled ($M_z$=10kNm, t=-0.6s, $\gamma$=13.1)**

# V. Conclusions and Future Work

## A. Conclusions

This research successfully demonstrated the application of Control Barrier-Value Function Quadratic Programming for vehicle lateral stability control, bridging the gap between theoretical safety guarantees and practical implementation. The CBVF framework unified Hamilton-Jacobi reachability analysis with control barrier function constraints, enabling systematic construction of safety controllers without manual barrier function design. The key contributions of this work include: (1) successful computation of controlled invariant sets for nonlinear vehicle dynamics incorporating tire saturation effects through HJ reachability analysis, (2) synthesis of a real-time implementable safety controller that maintains formal safety properties while producing smooth control signals, and (3) comprehensive validation demonstrating robust performance across varying surface conditions despite being tuned only for high-grip scenarios. The experimental results revealed compelling evidence for the practical viability of CBVF-based safety control. On low-grip surfaces, the safety filter achieved up to 82% reduction in tracking error and 75% reduction in control effort compared to baseline LQR control, successfully preventing vehicle spinout in all stable configurations. Critically, these improvements were achieved without surface-specific calibration, addressing a fundamental challenge in automotive control systems where traditional ESC requires extensive tuning across multiple operating conditions. The parameter study uncovered an unexpected dual nature of the CBVF controller. Configurations with constrained safe sets $M_{z_{\text{MAX}}} = 1$ kNm, $t_{\text{end}} = 0.2$ s operated as traditional boundary enforcing safety filters, while configurations with larger safe sets $M_{z_{\text{MAX}}} = 10$ kNm, $t_{\text{end}} = 0.6$ s evolved into gradient-informed augmentation systems that leveraged the value function's rich information content beyond the zero-level set. This latter configuration demonstrated superior robustness, maintaining stable operation through grip transitions that caused catastrophic failure in more conservative configurations.

The successful application of a 2-DOF bicycle model controller to a high-fidelity simulation validates the practical applicability of reduced-order safety control design, despite the absence of formal guarantees when crossing model complexity boundaries. This empirical validation contributes essential evidence to the ongoing theoretical development of reduced-order safety control methodologies.

## B. Future Work

Several research directions emerge from this work that could significantly advance the practical deployment and theoretical understanding of CBVF-based vehicle stability control.

Investigation of alternative vehicle state representations presents an immediate opportunity for enhanced practical implementation. Exploring state formulations based on yaw rate and lateral acceleration, rather than sideslip angle, would leverage sensors already present in production vehicles while potentially improving real-world applicability. However, this transition requires careful analysis of the information content trade-offs, as these alternative states may provide less direct insight into vehicle stability margins. Additionally, systematic evaluation of the CBVF controller's sensitivity to measurement noise and sensor uncertainties would establish critical implementation requirements for production deployment.

Development of realistic approximations for maximum control authority achievable through differential braking would eliminate the need for extensive parameter studies during controller design. The approach used requires empirical tuning of the theoretical $M_{z_{\text{MAX}}}$ parameter, which complicates practical implementation. Establishing physics-based relationships between brake system characteristics and achievable control moments would enable direct parameter specification and simplify the controller tuning.

A formal mathematical analysis of system reduction techniques for vehicle dynamics models could unlock powerful theoretical guarantees for stability and safety preservation. While this work demonstrates empirical success with bicycle model reduction, establishing rigorous conditions under which reduced-order models preserve the safety properties of high-dimensional systems would provide theoretical foundations for broader application. Such analysis could leverage recent developments in singular perturbation theory and geometric reduction methods to provide safety certificates for reduced-order controller design.

The rapidly evolving field of Hamilton-Jacobi reachability analysis offers significant opportunities for enhanced computational capabilities and system complexity handling. Advanced methods such as DeepReach neural network approximations, Koopman operator-based spectral approaches, and Hopf Hamilton-Jacobi formulations could enable direct computation on higher-dimensional vehicle models using only measurable states. These developments would eliminate model reduction requirements while potentially capturing coupling effects between lateral and longitudinal dynamics, suspension behavior, and aerodynamic influences that current simplified models neglect.

# References


[1] Ferguson, S. A., "The Effectiveness of Electronic Stability Control in Reducing Real-World Crashes: A Literature Review," *Traffic Injury Prevention*, Vol. 8, No. 4, 2007, pp. 329–338. https://doi.org/10.1080/15389580701588949.

[2] Cummings, P., McKnight, B., Rivara, F. P., and Grossman, D. C., "Association of driver air bags with driver fatality: a matched cohort study," *BMJ*, Vol. 324, No. 7346, 2002, pp. 1119–1122. https://doi.org/10.1136/bmj.324.7346.1119.

[3] Mueller, A., Achenbach, W., Schindler, E., Wohland, T., and Mohn, F. W., "Electronic stability program - the new active safety system of Mercedes-Benz," *Automobiltechnische Zeitschrift*, Vol. 96, No. 11, 1994.

[4] Zhu, Z., Tang, X., Qin, Y., Huang, Y., and Hashemi, E., "A Survey of Lateral Stability Criterion and Control Application for Autonomous Vehicles," Vol. 24, No. 10, 2023, pp. 10382–10399. https://doi.org/10.1109/TITS.2023.3280200.

[5] Aripin, M. K., Md Sam, Y., Danapalasingam, K. A., Peng, K., Hamzah, N., and Ismail, M. F., "A Review of Active Yaw Control System for Vehicle Handling and Stability Enhancement," Vol. 2014, No. 1, 2014, p. 437515. https://doi.org/10.1155/2014/437515.

[6] Pacejka, H. B., "Non-linearities in Road Vehicle Dynamics," Vol. 15, No. 5, 1986, pp. 237–254. https://doi.org/10.1080/00423118608968854.

[7] Rajamani, R., *Vehicle Dynamics and Control*, Mechanical Engineering Series, Springer US, 2012. https://doi.org/10.1007/978-1-4614-1433-9.

[8] Inagaki, S., "Analysis on vehicle stability in critical cornering using phase-plane method," 1994.

[9] Bobier-Tiu, C. G., Beal, C. E., Kegelman, J. C., Hindiyeh, R. Y., and Gerdes, J. C., "Vehicle control synthesis using phase portraits of planar dynamics," *Vehicle System Dynamics*, Vol. 57, 2019, pp. 1318–1337. https://doi.org/10.1080/00423114.2018.1502456.

[10] Németh, B., Gáspár, P., and Péni, T., "Nonlinear analysis of vehicle control actuations based on controlled invariant sets," Vol. 26, No. 1, 2016, pp. 31–43. https://doi.org/10.1515/amcs-2016-0003.

[11] Kumar, A., Umathe, B., Vaidya, U., and Kelkar, A., "Identifying lateral stability regions in vehicle dynamics: a Koopman spectral approach," Vol. 13, No. 1, 2025, p. 40. https://doi.org/10.1007/s40435-024-01562-w.

[12] Bansal, S., Chen, M., Herbert, S., and Tomlin, C. J., "Hamilton-Jacobi reachability: A brief overview and recent advances," *IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017, pp. 2242–2253. https://doi.org/10.1109/CDC.2017.8263977.

[13] Althoff, M., and Dolan, J. M., "Online Verification of Automated Road Vehicles Using Reachability Analysis," Vol. 30, No. 4, 2014, pp. 903–918. https://doi.org/10.1109/TRO.2014.2312453.

[14] Chen, M., and Tomlin, C. J., "Hamilton–Jacobi Reachability: Some Recent Theoretical Advances and Applications in Unmanned Airspace Management," Vol. 1, 2018, pp. 333–358. https://doi.org/10.1146/annurev-control-060117-104941.

[15] Ames, A. D., Coogan, S., Egerstedt, M., Notomista, G., Sreenath, K., and Tabuada, P., "Control Barrier Functions: Theory and Applications," *18th European Control Conference (ECC)*, 2019, pp. 3420–3431. https://doi.org/10.23919/ECC.2019.8796030.

[16] Ferraguti, F., Landi, C. T., Singletary, A., Lin, H.-C., Ames, A., Secchi, C., and Bonfè, M., "Safety and Efficiency in Robotics: The Control Barrier Functions Approach," Vol. 29, No. 3, 2022, pp. 139–151. https://doi.org/10.1109/MRA.2022.3174699.

[17] Ames, A. D., Xu, X., Grizzle, J. W., and Tabuada, P., "Control Barrier Function Based Quadratic Programs for Safety Critical Systems," *IEEE Transactions on Automatic Control*, Vol. 62, No. 8, 2017, pp. 3861–3876. https://doi.org/10.1109/TAC.2016.2638961.

[18] Wabersich, K. P., Taylor, A. J., Choi, J. J., Sreenath, K., Tomlin, C. J., Ames, A. D., and Zeilinger, M. N., "Data-Driven Safety Filters: Hamilton-Jacobi Reachability, Control Barrier Functions, and Predictive Methods for Uncertain Systems," Vol. 43, No. 5, 2023, pp. 137–177. https://doi.org/10.1109/MCS.2023.3291885.

[19] Hsu, K. C., Hu, H., and Fisac, J. F., "The Safety Filter: A Unified View of Safety-Critical Control in Autonomous Systems," Vol. 7, 2024, pp. 47–72. https://doi.org/10.1146/annurev-control-071723-102940.

[20] Tan, D. C. H., Acero, F., McCarthy, R., Kanoulas, D., and Li, Z., "Value Functions are Control Barrier Functions: Verification of Safe Policies using Control Theory," , December 2023. https://doi.org/10.48550/arXiv.2306.04026.

[21] Choi, J. J., Lee, D., Sreenath, K., Tomlin, C. J., and Herbert, S. L., "Robust Control Barrier–Value Functions for Safety-Critical Control," *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021, pp. 6814–6821. https://doi.org/10.1109/CDC45484.2021.9683085.

[22] Mitchell, I. M., "The Flexible, Extensible and Efficient Toolbox of Level Set Methods," Vol. 35, No. 2, 2008, pp. 300–329. https://doi.org/10.1007/s10915-007-9174-4.

[23] Li, T., and Jayawardhana, B., "Collision-Free Source Seeking Control Methods for Unicycle Robots," Vol. 70, No. 3, 2025, pp. 2020–2027. https://doi.org/10.1109/TAC.2024.3486654.

[24] Ganai, M., Gao, S., and Herbert, S. L., "Hamilton-Jacobi Reachability in Reinforcement Learning: A Survey," *IEEE Open Journal of Control Systems*, Vol. 3, 2024, pp. 310–324. https://doi.org/10.1109/OJCSYS.2024.3449138.

[25] Vandi, G., Moro, D., Ponti, F., Parenti, R., and Einaudi, G., "Vehicle Dynamics Modeling for Real-Time Simulation," *11th International Conference on Engines & Vehicles*, SAE International, 2013. https://doi.org/10.4271/2013-24-0144.

[26] Hirsch, D., and Herbert, S., "Approximate Hamilton-Jacobi Reachability Analysis for a Class of Two-Timescale Systems, with Application to Biological Models," *ArXiv*, 2025, p. arXiv:2503.11021v2.

[27] Righetti, G., Bertipaglia, A., Shyrokau, B., and Loenzo, B., "NMPC-Based Stability Control without Direct Sideslip Angle Regulation," *International Journal of Vehicle Mechanics and Mobility*, 2025.

[28] Sharpless, W., Shinde, N., Kim, M., Chow, Y. T., and Herbert, S., "Koopman-Hopf Hamilton-Jacobi Reachability and Control," , June 2025. https://doi.org/10.48550/arXiv.2303.11590.

[29] Pacejka, H., and Besselink, I. J. M., *Tire and Vehicle Dynamics*, Elsevier Science & Technology, 2012.

[30] Indu, K., and Kumar, M. A., "Electric Vehicle Control and Driving Safety Systems: A Review," *IETE Journal of Research*, Vol. 69, No. 1, 2023, pp. 482–498. https://doi.org/10.1080/03772063.2020.1830862.

[31] Wang, J.-X., Chen, N., Pi, D.-W., and Yin, G.-D., "Agent-based coordination framework for integrated vehicle chassis control," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, Vol. 223, No. 5, 2009, pp. 601–621. https://doi.org/10.1243/09544070JAUTO1015.

[32] Puscul, D., Lex, C., Vignati, M., and Shao, L., "A Literature Survey on Sideslip Angle Estimation Using Vehicle Dynamics Based Methods," *IEEE Access*, Vol. 12, 2024, pp. 70263–70277. https://doi.org/10.1109/ACCESS.2024.3402429.

[33] United Nations Economic Commission for Europe, "Regulation No. 140 – Uniform provisions concerning the approval of passenger cars with regard to Electronic Stability Control (ESC) systems," UNECE Regulation ECE/TRANS/WP.29/2014/59, United Nations Economic Commission for Europe, Geneva, Switzerland, 2014.

# Part II

## Additional Results

# 4

# Additional Testing

In the paper part of the thesis, results with an LQR as a base controller were presented. Additional tests were performed with an Industry sideslip controller to evaluate the CBVF safety filter's performance across different controller architectures. These results were not included in the main paper since the Industry controller violates the Lipschitz continuity assumption required for control barrier function theory. However, the cross-controller validation provides important insights into the practical limitations and broader applicability of the CBVF approach.

## 4.1. Industry Controller Performance Analysis

The Industry controller results demonstrate that the CBVF safety filter provides measurable benefits beyond simple baseline controllers, though with important limitations. The quantitative results are presented in Table 4.1 with corresponding time traces in Figure 4.1.

| Surface Condition | Industry | | Industry W/ SF | | Improvement | |
|---|---|---|---|---|---|---|
| | Tracking Error | Control Effort | Tracking Error | Control Effort | Tracking | Control |
| High Mu | 6.27 | 0.017 | 5.03 | 0.020 | 19.8% | -17.6% |
| Low Mu | 24.78 | 0.354 | 10.00 | 0.069 | 59.6% | 80.5% |
| Late Mu Transition | 25.79 | 0.323 | 11.00 | 0.151 | 57.3% | 53.3% |
| Early Mu Transition | 15.44 | 0.133 | 11.21 | 0.077 | 27.4% | 42.1% |

**Table 4.1:** Industry controller performance comparison across surface conditions

The Industry controller demonstrates superior baseline performance compared to LQR, particularly in challenging conditions. On high-grip surfaces, the unaugmented Industry controller achieves more precise tracking (6.27 vs 4.52 RMSE) though with slightly higher baseline error than LQR. However, fundamental limitations persist. The Industry controller still experiences complete spin-out on the low-grip surface, indicating that sophisticated threshold-based logic cannot overcome the challenges of operating beyond design parameters. The CBVF safety filter provides consistent improvements across all scenarios, though with smaller relative gains compared to LQR augmentation. The 59.6% improvement on low grip surfaces, while substantial, reflects the Industry controller's better baseline capability. More significantly, the 80.5% reduction in control effort demonstrates the safety filter's ability to prevent the oscillatory behavior characteristic of threshold-based controllers under extreme conditions.

The late transition scenario reveals critical limitations in the CBVF approach when integrated with sophisticated base controllers. Unlike other test cases where improvements exceed 50%, this scenario shows more modest gains (57.3%) with sideslip still exceeding 20°. Analysis of the moment traces identifies the root cause: the Industry controller intermittently disables both itself and the safety filter, as evidenced by abrupt transitions to zero applied moment. This override behavior fundamentally undermines the safety guarantees that CBF theory provides. When the base controller can unilaterally shut down the safety layer, the theoretical forward invariance properties no longer hold. The system effectively reverts to open-loop

dynamics during these periods, explaining why dangerous sideslip levels still occur despite safety filter operation. The control effort patterns reveal additional insights. The unaugmented Industry controller exhibits higher-frequency moment applications and erratic spikes during surface transitions, while the safety-augmented version maintains smoother control action when active. However, the override capability creates discontinuities that compromise overall system stability.

The Industry controller results highlight significant challenges for practical CBVF deployment. The consistent $\mu$ aware reference following behavior when the safety filter remains active demonstrates that the value function captures fundamental vehicle stability relationships rather than controller-specific artifacts. However, the override capability severely limits the ”one-shot” robustness claims established with the LQR baseline. While the Lipschitz continuity violation technically invalidates CBF’s theoretical guarantees, the practical performance demonstrates good behavior when the system remains active. This suggests the theoretical framework retains value even when implementation details deviate from idealized assumptions.

The results indicate that CBVF implementation requires careful consideration of base controller architecture and potentially redesigned integration approaches. Existing controllers with self-preservation logic designed to prevent actuator damage or system instability create fundamental conflicts with continuous safety filter operation. Future implementations must address these integration challenges to realize the full potential of the CBVF approach. Despite these limitations, the consistent performance improvements across both simple and sophisticated baseline controllers suggest that CBVF safety filtering addresses fundamental limitations in how controllers handle extreme nonlinear conditions. The approach shows promise for retrofit applications, but requires thoughtful integration design to prevent safety system bypass during critical events.

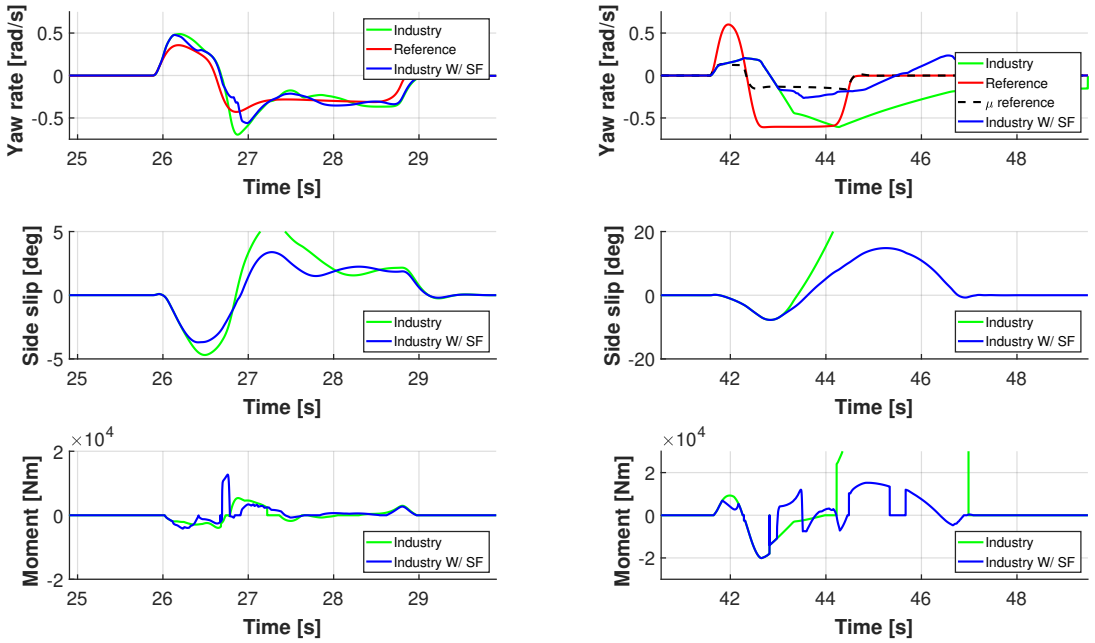


**(a)** High grip surface at 100 km/h




**(b)** Low grip surface at 50 km/h

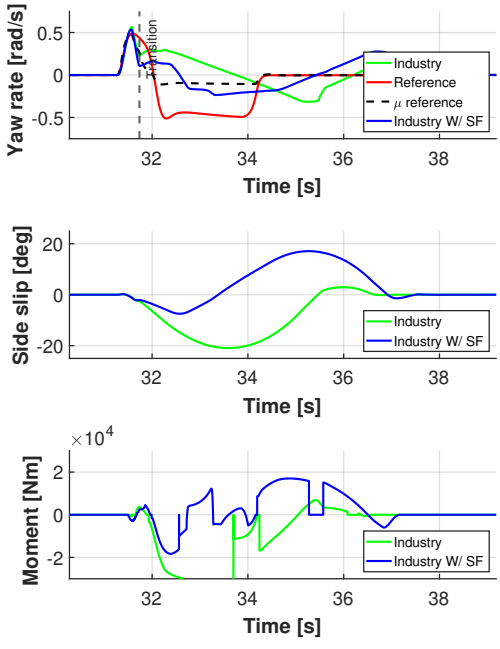


**(c)** High to low grip transition during initial turning at 70 km/h

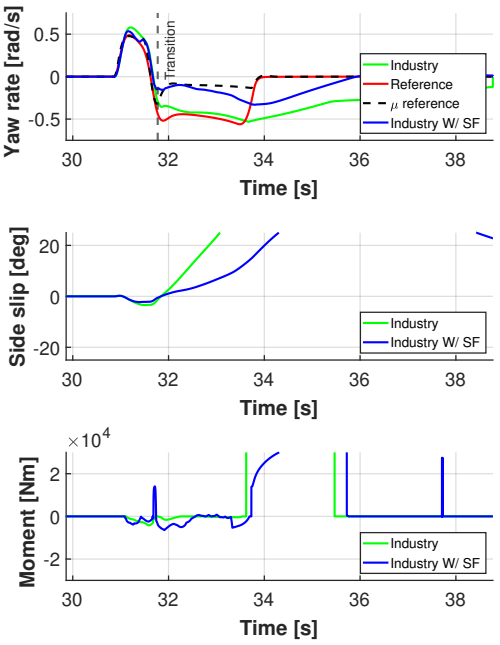


**(d)** High to low grip transition during dwell at 70 km/h

**Figure 4.1:** Comparison of Industry Controller with and without the CBVF safety controller enabled

# 5 Verification

This chapter presents additional work performed to ensure the correctness of the code used in this thesis. The verification procedure consists of three parts. First, the modified `hj_reachability`[1] package is tested against results from the reference paper [21]. Second, the reachability calculation results are verified against another technique for generating a safe phase plane for vehicle lateral dynamics. Finally, the CBVF formulation is tested to verify its ability to perform as a safety filter.

## 5.1. HJ Solver

To fully replicate the CBVF formulation, the `hj_reachability` package had to be extended with the time discount factor $\gamma$. To verify the implementation correctness, value functions for a simple problem were compared between the reference paper and the new implementation. Additional checks were performed against analytical results.

Starting with the comparison to previous work, the double integrator system used is shown in state-space form in Equation 5.1. The results were verified for $\gamma = 0$ and $\gamma = 0.5$. This ensured that setting $\gamma$ to zero allows the formulation to revert to a regular HJ PDE. The cross-check results are shown in Figure 5.1. Direct numerical comparisons were not possible since the original data is unavailable. However, examining the results shows that the modified package produces correct solutions. For $\gamma = 0$, the results are identical to the regular HJ solution. When $\gamma$ is set to 0.5, the change in solution matches that in the reference paper. There is a pronounced bulge of the function along the long side of the zero level set. The side profile also resembles a parabola in both solutions. Based on visual inspection, the two solvers produce the same solution.

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u \tag{5.1}$$

Another verification step involves checking if the solution follows the analytical solution. While the complete zero-level set does not have a closed-form analytical solution, parts of it do. The zero level set should consist of parabolic segments [22]. The analysis is presented in Figure 5.2. As indicated by the $R^2$ values, the top and bottom segments follow quadratic relations, confirming that the results adhere to the expected analytical solution.

The final step in verifying the HJ solver involves comparing the backward reachable set of a nonlinear bicycle model with another technique. Starting from the same grid points and using the same duration as the HJ reachability analysis, integrations of the nonlinear bicycle model were performed. The comparison results are shown in Figure 5.3. Only trajectories that reached the target set were plotted. As evident from the figure, the two methods for generating the safe phase-plane agree with one another, providing additional confidence in the correctness of the HJ solution.

[1] https://github.com/StanfordASL/hj_reachability

## 5.2. Quadratic Programming Controller

The quadratic programming controller verification was performed in two parts. First, results from the thesis implementation were compared to those from the reference paper. Figure 5.4 and Figure 5.5 show these comparisons. The results are not identical, which stems from different tuning of the underlying PI controllers. The reference paper did not specify the gains used, so this represents the best approximation based on the resulting trajectories. Nevertheless, the behaviors are similar, indicating a correct implementation.

For $\gamma = 0$, the controller refuses to reduce the CBVF value, resulting in a trajectory ending on the zero velocity line. The control inputs between the two controllers show similar trends. For $\gamma = 0.5$, the results also confirm correct implementation. The controller allows decay in the value function, and the trajectories mimic each other. Again, the matching is not exact, but this is due to different controller gains.

The final verification step for the QP controller involved checking if it can intervene appropriately when applied to the nonlinear bicycle model. This tests the entire pipeline, from the HJ solution to the final QP controller implementation. The safety test results are presented in Figure 5.6. The test initialized the system at 0 rad/s and 0.2 rad of sideslip. The underlying controller outputs a constant input of 10 kNm.

In Figure 5.6a, two simulation traces are visible: one with the safety controller enabled and one without. The results confirm that the safety controller can perform its function effectively. While the system without safety control quickly leaves the BRS boundary, the system with the safety controller maintains itself on the BRS boundary. From Figure 5.6b, it is evident that the controller intervenes minimally. Examining the CBVF value function shows it decays rapidly toward zero but never crosses it. This is exactly the expected behavior. Based on all the testing above, it is clear that the CBVF framework has been implemented correctly.

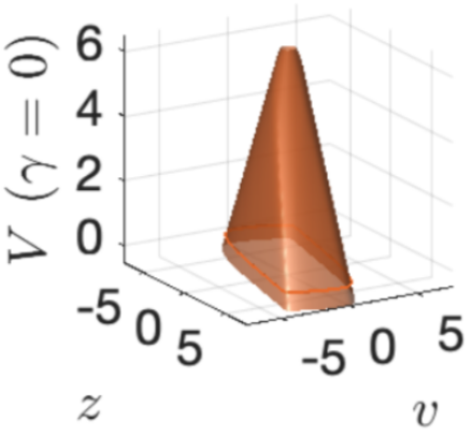


**(a)** Reference result for $\gamma = 0$

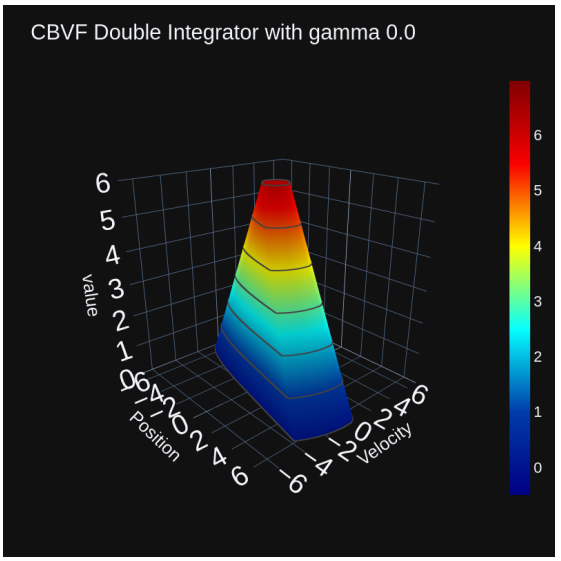


**(b)** Modified `hj_reachability` result for $\gamma = 0$

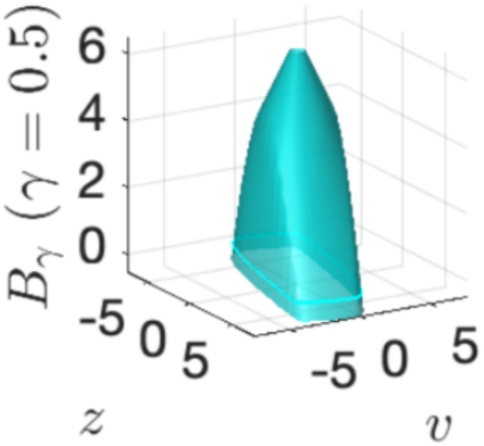


**(c)** Reference result for $\gamma = 0.5$

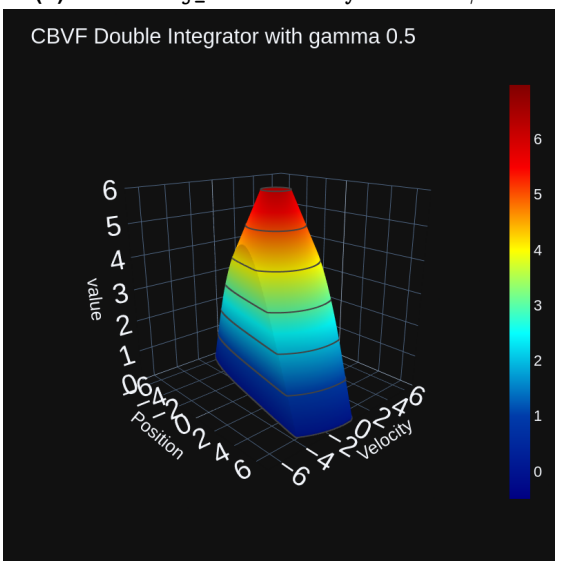


**(d)** Modified `hj_reachability` result for $\gamma = 0.5$

**Figure 5.1:** Cross check of value function results between modified `hj_reachability` and reference paper

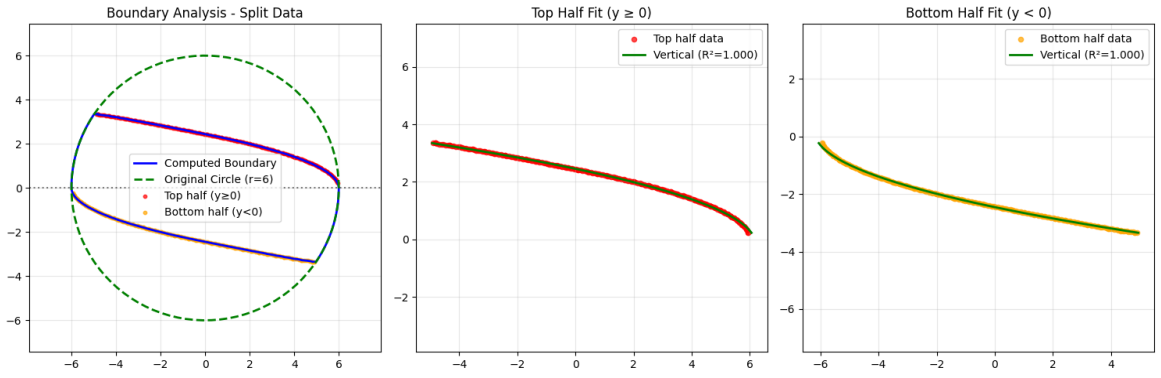


**Figure 5.2:** Results of quadratic fit on the zero level set of double integrator

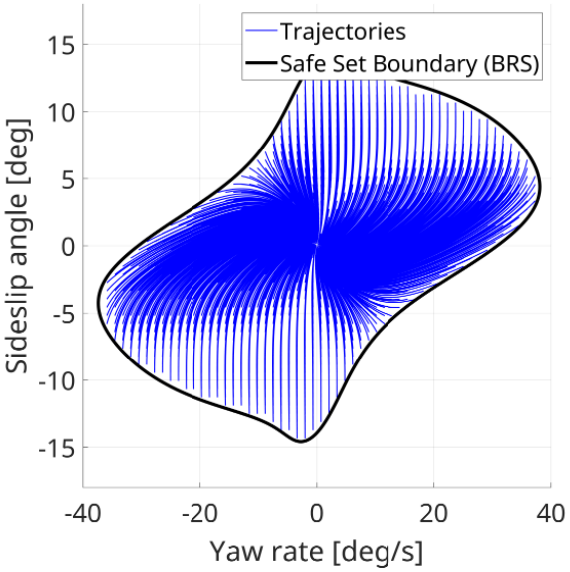


**Figure 5.3:** Comparison between HJ zero level set and time integration starting from many grid points

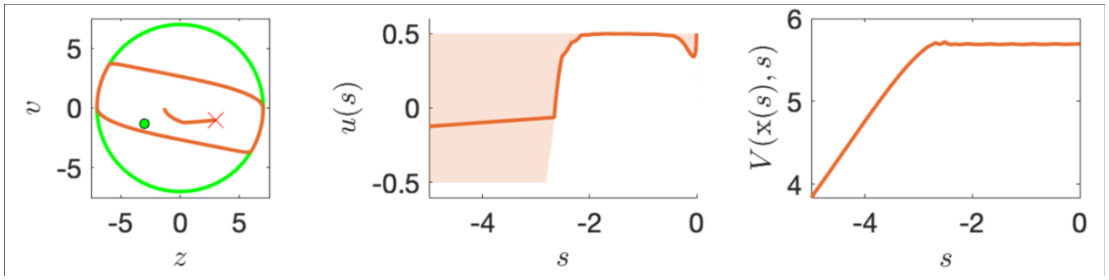


**(a)** Results from the reference paper

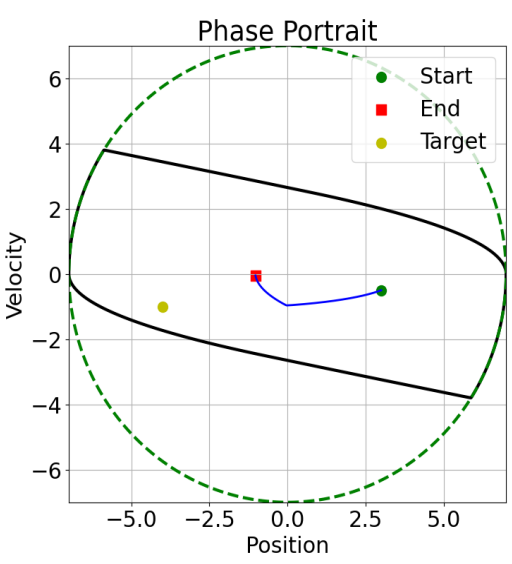


**(b)** Trajectory for the controller used in the thesis

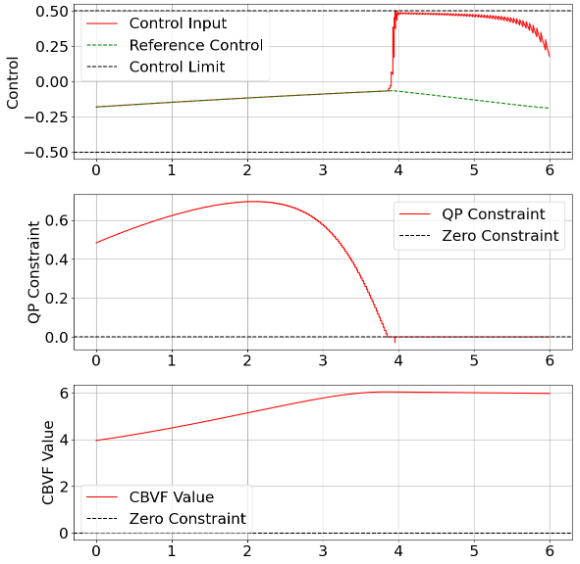


**(c)** Time traces for the controller used in the thesis

**Figure 5.4:** Quadratic Controller verification for $\gamma = 0$

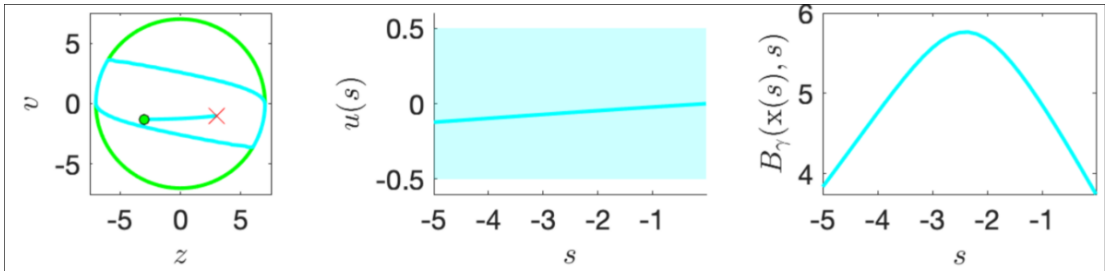


**(a)** Results from the reference paper

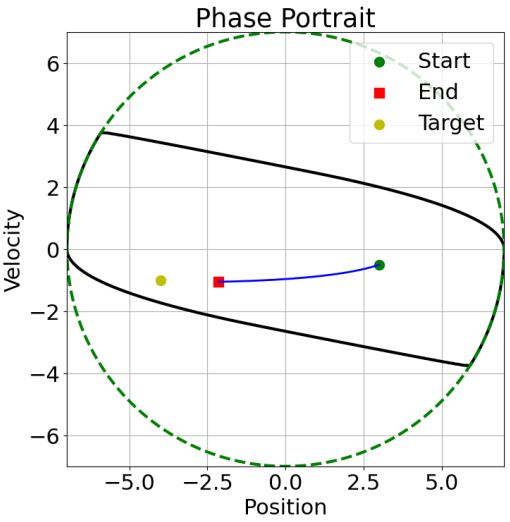


**(b)** Trajectory for the controller used in the thesis

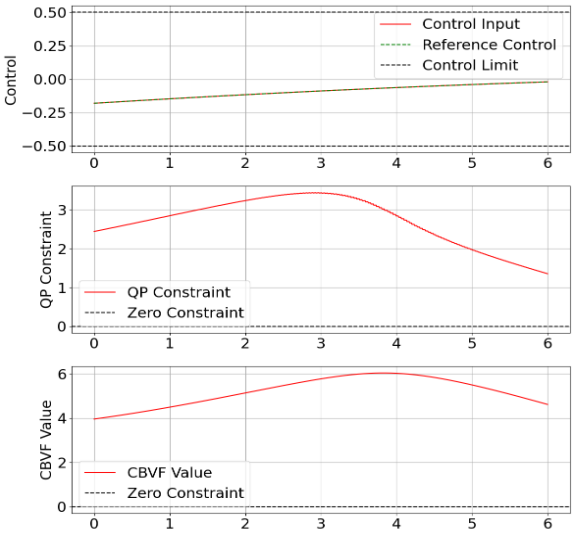


**(c)** Time traces for the controller used in the thesis

**Figure 5.5:** Quadratic Controller verification for $\gamma = 0.5$

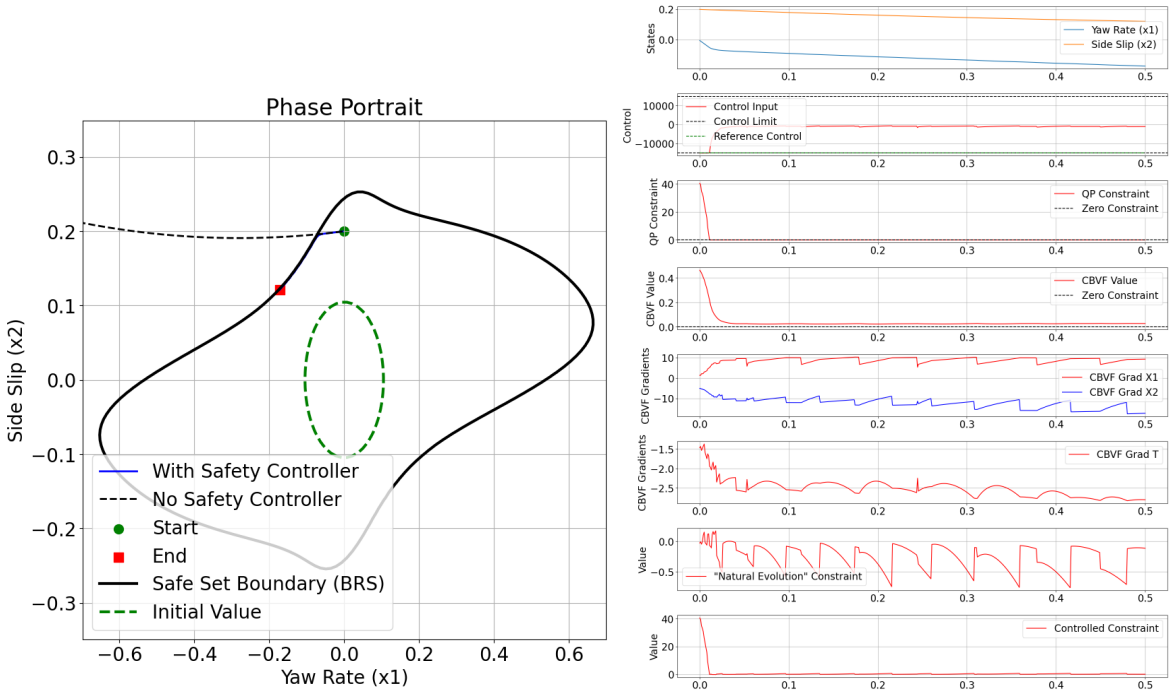


**(a)** Trajectories with and without the CBVF safety controller present

**(b)** Time traces of the CBVF safety controller enabled test

**Figure 5.6:** Verification on nonlinear bicycle model

# 6

# Conclusion

## 6.1. Closing Remarks

This thesis successfully demonstrated the practical viability of Control Barrier-Value Function Quadratic Programming (CBVF-QP) for vehicle lateral stability control, bridging the gap between theoretical safety guarantees and real-world automotive applications. The research established that Hamilton-Jacobi reachability analysis, when unified with control barrier function frameworks, provides a systematic approach for ensuring vehicle stability without requiring surface-specific calibration or manual barrier function design.

The key achievement of this work lies in demonstrating that a safety controller computed on a simplified nonlinear bicycle model can effectively stabilize a high-fidelity vehicle simulation across dramatically different operating conditions. Despite the significant model mismatch between the 3-DOF bicycle model used for BRT computation and the complex CarMaker simulation platform, the CBVF controller maintained vehicle stability across all tested scenarios, including extreme transitions from high to low grip surfaces ($\mu = 1.0$ to $\mu = 0.2$).

The parameter study revealed fundamental insights into how different BRT design parameters affect controller behavior. The distinction between narrow and wide BRT configurations represents two complementary control philosophies: boundary-focused safety enforcement versus gradient-based steady-state tracking. This flexibility, achieved through tuning of the relaxation parameter $\gamma$, maximum control authority, and prediction horizon, enables adaptation to specific vehicle characteristics without redesigning the underlying safety framework.

Perhaps most significantly, the “one-shot” capability demonstrated throughout the testing eliminates a critical limitation of current stability control systems. The controller’s ability to maintain performance margins while operating on surfaces with 80% less available grip than its tuning conditions represents a substantial advancement over threshold-based ESC systems that require extensive calibration for different road conditions. This parameter error resilience suggests that CBVF-based safety filters could reduce development complexity while improving safety margins in production vehicles.

## 6.2. Research Questions

The research questions posed in Section 2.2 are addressed below with the findings from this thesis:

### Main Research Question

*How can Hamilton-Jacobi reachability analysis be effectively applied to compute safe operating envelopes of a non-linear vehicle bicycle model and subsequently used in creating an effective safety controller?*

This thesis definitively establishes that HJ reachability analysis can be successfully applied to vehicle lateral stability through the CBVF framework. The approach computes maximal safe operating envelopes that capture nonlinear tire saturation effects while enabling real-time safety intervention through quadratic programming. The resulting controller achieved up to 78% improvement in tracking performance and 62% reduction in control effort compared to baseline controllers, while preventing vehicle spin-out in scenarios where traditional controllers failed.

### Sub-Question 1: Computation of Safe Sets

*How can Hamilton-Jacobi reachability analysis be applied to compute accurate safe operating envelopes for the nonlinear bicycle model under different velocity and road friction conditions?*

The research demonstrated that HJ reachability effectively computes controlled invariant sets for the nonlinear bicycle model with nonlinear tire dynamics. The backward reachable tube computation successfully captured the complex relationship between yaw rate and sideslip angle across different operating conditions. The choice of steady-state corrected target sets, rather than zero-based targets, proved crucial for accounting for the effects of steerable front wheels. The 0.2–0.6 second time horizons provided a practical balance between safety guarantees and operational freedom, with longer horizons enabling more anticipatory control actions.

### Sub-Question 2: Interpolation Between Safe Sets

*What methods can be used to accurately interpolate between precomputed safe sets?*

While formal guarantees for interpolated safe sets remain an open theoretical challenge, the practical implementation demonstrated that direct application of precomputed safe sets across different operating conditions was surprisingly effective. The value function's implicit encoding of vehicle dynamics provided sufficient information for the controller to adapt to varying conditions without explicit interpolation. The controller's successful operation across the tested $\mu$ range (0.2 to 1.0) suggests that the HJ value function captures fundamental stability relationships that transcend specific parameter values. This suggests that there might not be a need for interpolating between HJ value functions. The relaxation factor in the CBVF formulation and the possibility of including disturbances in the HJ calculations might be enough.

### Sub-Question 3: Safety Controller Design

*How can the computed and interpolated safe sets be effectively utilized in a safety controller for vehicle lateral stability?*

The CBVF-QP formulation proved highly effective for utilizing HJ-derived safe sets. The quadratic program implementation with relaxation parameter $\gamma$ enabled smooth control interventions that minimally modified nominal control inputs. The controller architecture successfully balanced multiple objectives: maintaining safety through forward invariance constraints, preserving performance through minimal intervention, and ensuring implementability through bounded control actions. The distinction between narrow and wide BRT behaviors revealed through parameter exploration provides designers with systematic tools for achieving desired control characteristics.

### Sub-Question 4: Validation of Safety Controller

*How can the developed safety controller be checked to be applicable in real-world scenarios?*

High-fidelity validation using CarMaker with a validated passenger model confirmed the controller's practical applicability despite the simplified model used for its design. The controller passed all ECE R140 regulatory requirements while demonstrating robust performance across varying surface conditions. The successful prevention of spin-out scenarios and maintenance of stability during friction transitions validates the approach for safety-critical applications. However, testing with the Industry controller revealed integration challenges with existing control architectures that include override logic, highlighting the need for careful system-level design in production implementations.

## 6.3. Scientific Contributions

This thesis makes several significant contributions to the field of vehicle stability control:

1. **First systematic application of CBVF to vehicle lateral stability**: This work represents the first comprehensive implementation of Hamilton-Jacobi reachability-based safety filtering for vehicle lateral dynamics, establishing a new paradigm for proactive stability control.
2. **Demonstration of model reduction validity**: The successful stabilization using controllers designed on simplified models provides empirical evidence that appropriate model reduction preserves essential safety properties, addressing a critical gap in the theoretical understanding of safety-critical control system design.

3. **Parameter error resilience without adaptation**: The controller's ability to maintain stability across an 80% reduction in available grip without retuning challenges conventional wisdom about the need for surface-specific calibration in vehicle stability systems.
4. **Unified framework bridging theory and practice**: By successfully implementing CBVF on a realistic vehicle model and validating against industry standards, this work bridges the gap between academic safety filter research and practical automotive applications.

# 7

# Recommendations

This chapter provides a brief overview of the primary recommendations for the future continuation of this research project.

**Recommendation 1: Develop Alternative State Representations Using Measurable Quantities**
Future research should prioritize developing CBVF formulations based on directly measurable vehicle states such as yaw rate and lateral acceleration, rather than sideslip angle. This transition would eliminate the need for complex state estimation in production vehicles while maintaining safety guarantees. Investigation should focus on whether the $\psi$-$a_y$ phase plane can provide equivalent safety characterization to the current $\beta$-$\psi$ representation, potentially leveraging the extensive sensor infrastructure already present in modern vehicles.

**Recommendation 2: Establish Formal Guarantees for Model Reduction Techniques**
A rigorous mathematical framework should be developed to bound the degradation of safety properties under model mismatch. This work should leverage recent advances in singular perturbation theory and geometric reduction methods to provide formal certificates that guarantee safety preservation when controllers designed on reduced-order models are applied to high-dimensional systems. Specific attention should be given to establishing conditions under which the bicycle model reduction preserves the controlled invariance properties of full vehicle dynamics.

**Recommendation 3: Redesign Control Architectures for Seamless Safety Filter Integration**
New control system architectures must be developed that fundamentally integrate safety filtering at the system level, preventing the override behaviors observed with existing Industry controllers. This requires rethinking the hierarchical structure of vehicle stability systems to ensure that safety constraints cannot be bypassed by lower-level self-preservation logic. The architecture should maintain Lipschitz continuity requirements while accommodating practical considerations such as actuator protection and fault handling.

**Recommendation 4: Extend Framework to Combined Lateral-Longitudinal Dynamics**
The CBVF framework should be extended to handle coupled lateral-longitudinal maneuvers and actuator dynamics. This extension should investigate whether recent developments in HJ reachability, such as DeepReach neural network approximations or Hopf formulations, can enable computation on higher-dimensional models without sacrificing real-time implementability. Priority should be given to capturing the coupling effects that become significant during combined braking and steering interventions.

**Recommendation 5: Develop Physics-Based Relationships for Control Authority Specification**
Systematic methods should be established for determining the theoretical $M_{z,\text{MAX}}$ parameter based on physical brake system characteristics. This would eliminate the current need for empirical tuning by creating direct mappings between brake pressure capabilities, tire-road friction, and achievable control moments. The approach should account for dynamic load transfer and provide conservative bounds suitable for safety-critical applications.

**Recommendation 6: Implement Advanced Computational Methods for Higher-Dimensional Systems**
Future work should explore advanced computational techniques including Koopman operator-based spectral approaches and GPU-accelerated local update schemes to enable direct BRT computation on models with more than 10 degrees of freedom. This would allow inclusion of suspension dynamics, tire temperature effects, and other phenomena currently neglected in the simplified bicycle model, potentially improving controller performance in extreme maneuvers.

# References


[1] K. C. Hsu et al. “The Safety Filter: A Unified View of Safety-Critical Control in Autonomous Systems”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 7 (July 2024). Publisher: Annual Reviews, pp. 47–72. DOI: 10.1146/annurev-control-071723-102940.

[2] S. Bansal et al. “Hamilton-Jacobi reachability: A brief overview and recent advances”. In: *IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE 56th Annual Conference on Decision and Control (CDC). Dec. 2017, pp. 2242–2253. DOI: 10.1109/CDC.2017.8263977.

[3] A. D. Ames et al. “Control Barrier Functions: Theory and Applications”. In: *2019 18th European Control Conference (ECC)*. 2019 18th European Control Conference (ECC). June 2019, pp. 3420–3431. DOI: 10.23919/ECC.2019.8796030.

[4] K. P. Wabersich et al. “Data-Driven Safety Filters: Hamilton-Jacobi Reachability, Control Barrier Functions, and Predictive Methods for Uncertain Systems”. In: *IEEE Control Systems Magazine* 43.5 (Oct. 2023). Conference Name: IEEE Control Systems Magazine, pp. 137–177. DOI: 10.1109/MCS.2023.3291885.

[5] Z. Zhu et al. “A Survey of Lateral Stability Criterion and Control Application for Autonomous Vehicles”. In: *IEEE Transactions on Intelligent Transportation Systems* 24.10 (Oct. 2023). Conference Name: IEEE Transactions on Intelligent Transportation Systems, pp. 10382–10399. DOI: 10.1109/TITS.2023.3280200.

[6] B. Németh et al. “Nonlinear analysis of vehicle control actuations based on controlled invariant sets”. In: *International Journal of Applied Mathematics and Computer Science* 26.1 (Mar. 2016), pp. 31–43. DOI: 10.1515/amcs-2016-0003.

[7] A. Kumar et al. “Identifying lateral stability regions in vehicle dynamics: a Koopman spectral approach”. In: *International Journal of Dynamics and Control* 13.1 (Jan. 2025), p. 40. DOI: 10.1007/s40435-024-01562-w.

[8] M. Althoff et al. “Online Verification of Automated Road Vehicles Using Reachability Analysis”. In: *IEEE Transactions on Robotics* 30.4 (Aug. 2014). Conference Name: IEEE Transactions on Robotics, pp. 903–918. DOI: 10.1109/TRO.2014.2312453.

[9] M. Chen et al. “Hamilton–Jacobi Reachability: Some Recent Theoretical Advances and Applications in Unmanned Airspace Management”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 1 (May 2018). Publisher: Annual Reviews, pp. 333–358. DOI: 10.1146/annurev-control-060117-104941.

[10] J. F. Fisac et al. “A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems”. In: *IEEE Transactions on Automatic Control* 64.7 (July 2019). Conference Name: IEEE Transactions on Automatic Control, pp. 2737–2752. DOI: 10.1109/TAC.2018.2876389.

[11] Y. Zhang et al. “Database Building and Interpolation for an Online Safe Flight Envelope Prediction System”. In: *Journal of Guidance, Control, and Dynamics* 42.5 (2019), pp. 1166–1174. DOI: 10.2514/1.G003834.

[12] M. Chen et al. “FaSTrack:A Modular Framework for Real-Time Motion Planning and Guaranteed Safe Tracking”. In: *IEEE Transactions on Automatic Control* 66.12 (Dec. 2021). Conference Name: IEEE Transactions on Automatic Control, pp. 5861–5876. DOI: 10.1109/TAC.2021.3059838.

[13] L. K. Chung et al. *Goal-Reaching Trajectory Design Near Danger with Piecewise Affine Reach-avoid Computation*. May 2024. DOI: 10.48550/arXiv.2402.15604.

[14] H. K. Sachs et al. “Automobile Stability – A Study of the Domain of Attraction”. In: *Vehicle System Dynamics* 6.2 (Sept. 1977), pp. 169–177. DOI: 10.1080/00423117708968531.

[15] Y. Shibahata et al. “Improvement of Vehicle Maneuverability by Direct Yaw Moment Control”. In: *Vehicle System Dynamics* 22.5 (Jan. 1993), pp. 465–481. DOI: 10.1080/00423119308969044.

[16] S. Inagaki. “Analysis on vehicle stability in critical cornering using phase-plane method”. In: *Proc.of AVEC ’94* (1994). Publisher: JSAE.

[17] H. B. Pacejka. “Non-linearities in Road Vehicle Dynamics”. In: *Vehicle System Dynamics* 15.5 (Jan. 1986), pp. 237–254. DOI: 10.1080/00423118608968854.

[18] R. Rajamani. *Vehicle Dynamics and Control*. Mechanical Engineering Series. Boston, MA: Springer US, 2012. DOI: 10.1007/978-1-4614-1433-9.

[19] X. Chen et al. “Reachability Analysis for Cyber-Physical Systems: Are We There Yet?” In: *NASA Formal Methods*. Ed. by J. V. Deshmukh et al. Vol. 13260. Series Title: Lecture Notes in Computer Science. Cham: Springer International Publishing, 2022, pp. 109–130. DOI: 10.1007/978-3-031-06773-0_6.

[20] I. M. Mitchell. “The Flexible, Extensible and Efficient Toolbox of Level Set Methods”. In: *Journal of Scientific Computing* 35.2 (June 2008), pp. 300–329. DOI: 10.1007/s10915-007-9174-4.

[21] J. J. Choi et al. “Robust Control Barrier–Value Functions for Safety-Critical Control”. In: *2021 60th IEEE Conference on Decision and Control (CDC)*. 2021 60th IEEE Conference on Decision and Control (CDC). ISSN: 2576-2370. Dec. 2021, pp. 6814–6821. DOI: 10.1109/CDC45484.2021.9683085.

[22] D. Liberzon. *Calculus of variations and optimal control theory : a concise introduction*. 1 online resource (xv, 235 pages) : illustrations vols. Princeton, New Jersey: Princeton University Press, 2012.