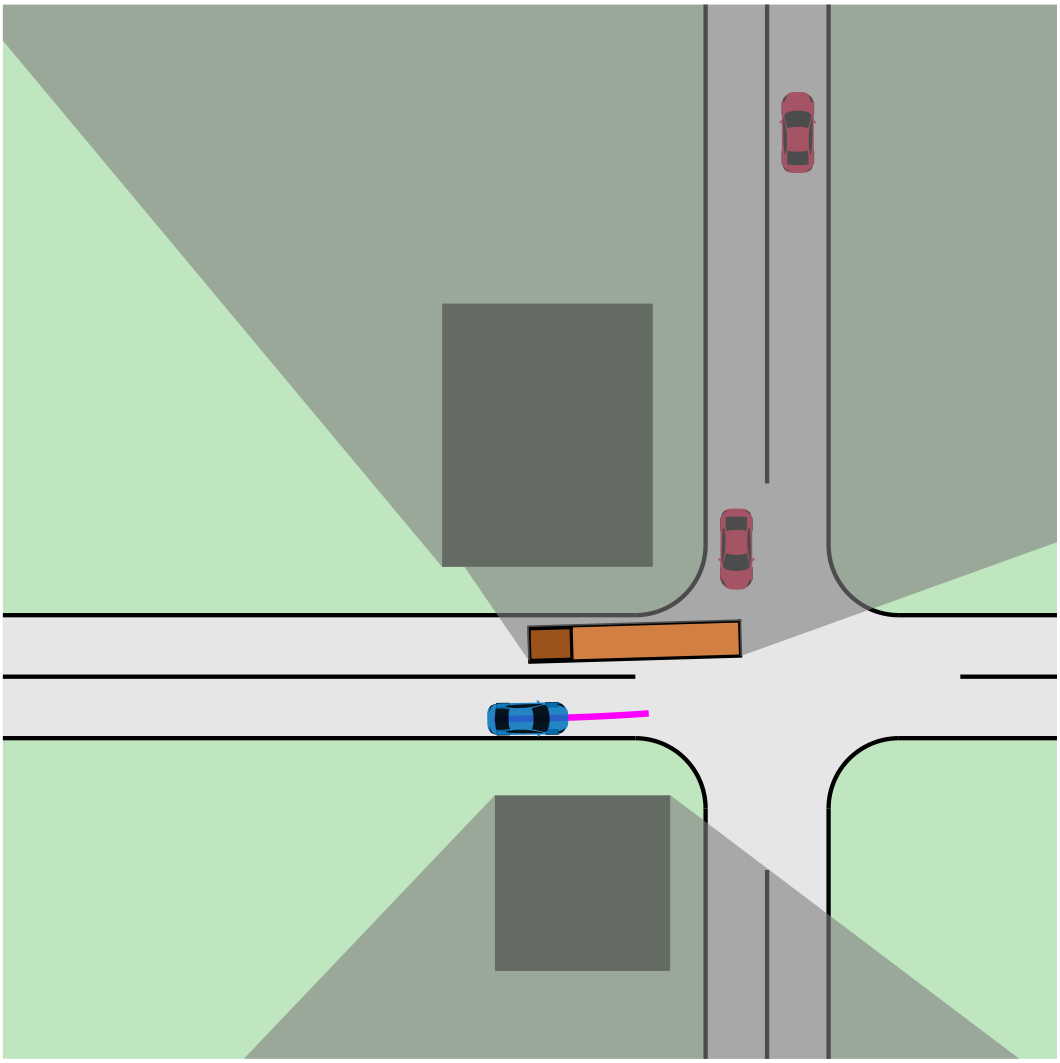

Robust MPC-based Trajectory Planning for Autonomous Driving under Occlusion

Junyan Tian & Abishek Swaminathan

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2025
www.chalmers.se

Master's thesis 2025

Robust MPC-based Trajectory Planning for Autonomous Driving under Occlusion

Junyan Tian

Abishek Swaminathan

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
Division of Systems, Control and Mechatronics
Chalmers University of Technology
Gothenburg, Sweden 2025

Robust MPC-based Trajectory Planning for Autonomous Driving under Occlusion
Junyan Tian, Abishek Swaminathan

Supervisor: Cheerudeep Chintha, Qualcomm
Supervisor & Examiner: Nikolce Murgovski, Department of Electrical Engineering

Master's Thesis 2025
Department of Electrical Engineering
Systems and Control
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

The authors would like to acknowledge that parts of this manuscript were edited with the assistance of a large language model to improve clarity and readability.

Cover: Autonomous vehicle trajectory planning under mixed occlusion scenarios. Blue ego vehicle with magenta planned trajectory navigates intersection occlusions caused by static infrastructure and a dynamic orange truck obscuring red cross-traffic vehicles.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2025

Robust MPC-based Trajectory Planning for Autonomous Driving under Occlusion
Junyan Tian & Abishek Swaminathan
Department of Electrical Engineering
Chalmers University of Technology

Abstract

With the rapid advancement and widespread adoption of autonomous driving, trajectory planning must not only ensure efficiency and comfort but also guarantee safety. Planning under occlusions, however, remains a longstanding challenge. In situations such as sudden pedestrian emergence, conventional methods based on sampling or optimization primarily account for visible obstacles, while end-to-end deep learning approaches, although capable of implicitly considering occlusion effects, often suffer from black-box characteristics and limited interpretability. To address this issue, we propose a lightweight occlusion-aware robust MPC trajectory planning module. The module can be seamlessly integrated into existing planners and is selectively activated in high-risk occlusion scenarios to enhance safety. Using reachable set analysis, we explicitly model hidden road users and their motion predictions, which are incorporated into a carefully designed Robust Model Predictive Control (RMPC) framework. Two representative cases, hidden pedestrians and hidden vehicles, are investigated through extensive simulation studies. Compared with a baseline MPC, our approach reduces the collision rate from over 10% to 0%, demonstrating its effectiveness in ensuring safe navigation under occlusions

Keywords: Trajectory Planning, Robust Model Predictive Control, Occlusion, Reachability Analysis, Collision Avodiance

Acknowledgements

We extend our sincere gratitude to Hoda Yarmohamadi Fredriksson for providing us with the invaluable opportunity to conduct our Master's thesis at Qualcomm. This experience has been truly enriching, offering us direct exposure to a high-technology company and allowing us to pursue our long-standing passion for autonomous driving research.

Our profound appreciation goes to our academic supervisor and examiner, Nikolce Murgovski, from Chalmers University of Technology. His insightful guidance, continuous support, and constructive feedback have been instrumental in shaping this project from its inception to its completion.

We are also deeply thankful to our industrial supervisor, Cheerudeep Chintha, for his unwavering support and expert guidance throughout the ideation phase and the execution of this project. His practical insights were invaluable. Furthermore, we wish to express our gratitude to Lizi Teng and Karl Rundstedt at Qualcomm for their insightful discussions and contributions, which significantly broadened our understanding of industry practices and challenges.

Finally, we are immensely grateful to Qualcomm for hosting us and providing a collaborative environment filled with wonderful colleagues, making this thesis journey a truly rewarding experience. Last but not least, we wholeheartedly thank our families for their unwavering support and encouragement, which were essential for the successful completion of this work.

Junyan Tian & Abishek Swaminathan, Gothenburg, September 2025

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

AD	Autonomous Driving
ADAS	Advanced Driver Assistance Systems
VLA	Vision-Language-Action
BRS	Backward Reachable Set
DP	Dynamic Programming
FRS	Forward Reachable Set
IPOPT	Interior Point OPTimizer
MILP	Mixed-Integer Linear Programming
MPC	Model Predictive Controller
NLP	Non Linear Programming
OCP	Optimal Control Problem
ODE	Ordinary Differential Equations
RK4	Fourth Order Runge-Kutta Method
RMPC	Robust Model Predictive Controller
VLA	Vision-Language-Action
VRU	Vulnerable Road User
RMS	Root Mean Square

Nomenclature

Below is the nomenclature of indices, sets, parameters, and variables that have been used throughout this thesis.

- Vectors are bold lower case letters "$\mathbf{x}$"
- Sets are in blackboard-bold letters "$\mathbb{S}$"
- Time derivatives are marked by "$\dot{x}$"

Sets

$\emptyset$	Empty set
$\mathbb{R}$	Set of Real Numbers
$\mathbb{U}$	Set of inputs
$\mathbb{X}_0$	Initial set
$\mathbb{T}$	Target set
$\mathbb{B}$	Backward reachable set
$\mathbb{O}_{\text{occ}}$	Geometrically characterized occlusion
$\mathbb{O}_{\text{obs}}$	Currently observable area
$\partial\mathbb{O}_{occ}$	Boundary of the visibility polygon
$\mathbb{L}$	Centerlines of traffic lanes
$\mathbb{N}$	Intersection nodes
$\mathbb{N}_{high}$	High risk node set
$\mathbb{N}_{low}$	Low risk node set
$\mathbb{N}^{near}$	Near node set
$\mathbb{N}^{far}$	Far node set

Parameters

L	Wheelbase of the ego vehicle
tf	Final time
N	Horizon
dt	Time interval

$v_{x,\max}^{ped}$	Maximum lateral velocity of pedestrian
$v_{y,\max}^{ped}$	Maximum longitudinal velocity of pedestrian
$v_{\max}^{veh}$	Maximum velocity of hidden vehicle
$v_{\min}^{veh}$	Minimum velocity of hidden vehicle
d_{safety}	Safety margin
L_{veh}	Length of a vehicle
W_{veh}	Width of a vehicle
W_{lane}	Width of a traffic lane

Variables

s	Position of the rear axle of the ego vehicle along the X axis
y	Position of the rear axle of the ego vehicle along the Y axis
θ	Heading Angle
δ	Steering Angle
v	Velocity
ω	Steering Rate
a	Acceleration
j	Jerk
θ_e	Heading error
λ	Normal Vector
$\alpha,\ \beta$	Shape Parameters for the beta distribution
J	Cost function
S	Slack
x_{ref}	Reference trajectory state vector
u_{ref}	Reference control input vector
$\mathbf{p}^{\text{near}}$	Near node of occluded lane segment for hidden vehicle
$\mathbf{p}^{\text{far}}$	Far node of occluded lane segment for hidden vehicle
Δt	Discrete time step for dynamic propagation of initial sets
ℓ_x	State tracking cost in RMPC stage cost
ℓ_u	Control effort cost in RMPC stage cost
ℓ_{terminal}	Terminal cost in RMPC objective function
ℓ_{slack}	Slack penalty for constraint relaxation in RMPC
Q	Weighting matrix for state tracking cost in RMPC

R	Weighting matrix for input effort in RMPC

Functions

$\mathcal{R}_F(\cdot)$	Forward reachable set calculation function
$\mathcal{R}_B(\cdot)$	Backward reachable set calculation function

Contents

List of Figures

List of Tables

1

Introduction

This chapter introduces the motivation, research problem, and proposed solution of the thesis. It begins by highlighting the challenge of sensor occlusions in autonomous driving, reviews the limitations of existing approaches, and formulates the research question. The chapter then presents the proposed robust planning framework and summarizes the key contributions before outlining the structure of the remainder of the document.

1.1 Background

Autonomous driving imposes stringent safety requirements, particularly in the presence of unforeseen environmental dynamics. Traditional modular trajectory planning pipelines [1], typically structured from perception to decision-making to planning, primarily focus on explicitly accounting for perceived, visible obstacles within the immediate environment. However, such pipelines inherently struggle to handle potential hazards emerging from unseen occluded regions, which represent critical safety blind spots.

End-to-end planners [2] and Vision-Language-Action (VLA) approaches [3], while potentially capable of implicitly correlating complex sensor inputs with occluded threats through deep learning, suffer from a black-box nature and limited interpretability. This lack of transparency is a major concern in safety-critical applications, where verifiability and predictability are essential.

Given the significant safety implications in high-occlusion, high-risk scenarios, an explicit, interpretable model-based approach is essential as a robust safety fallback. To this end, we introduce a Risk-Aware Replanning Module within our system architecture as shown in Figure 1.1. This module dynamically assesses occlusion risk using perception data and triggers a specialized risk-aware replanning process when a sufficiently high level of risk is detected. In such cases, it generates a cautious trajectory that proactively mitigates potential collisions. Under conditions of lower assessed risk, the system continues to execute the nominal trajectory provided by the primary planner.

Several strategies can be employed to incorporate occlusion awareness into trajectory planning, among which optimization-based methods are particularly effective.

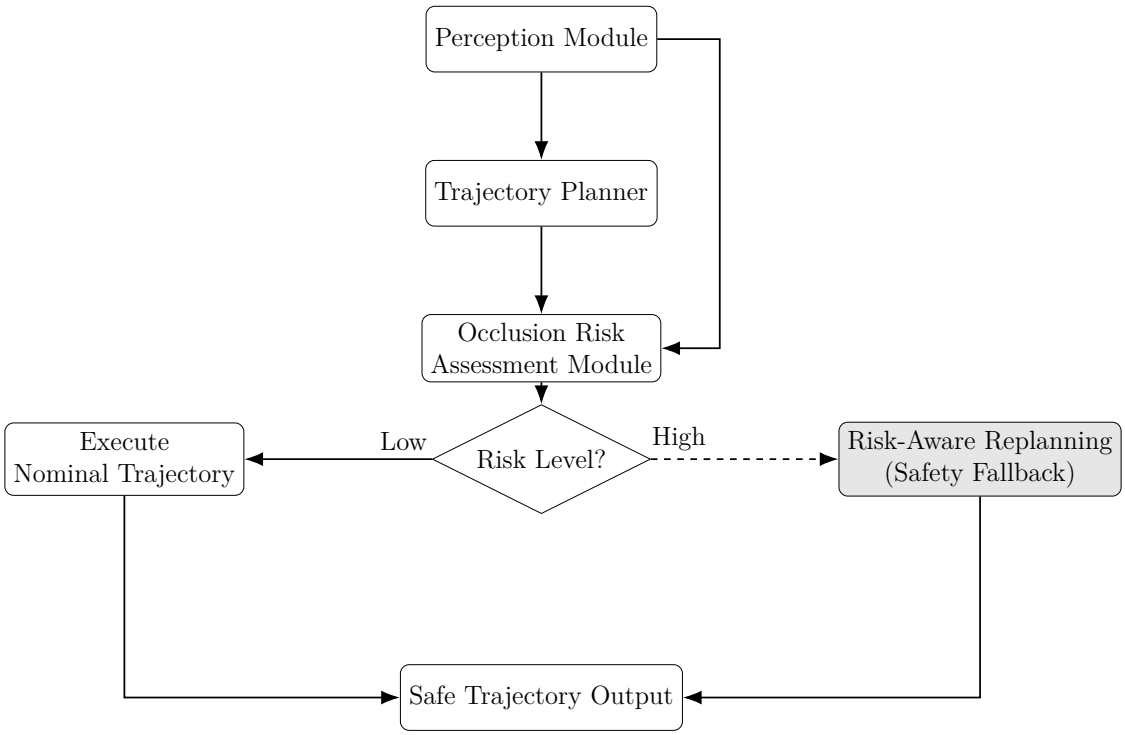

Figure 1.1: System architecture highlighting the occlusion-aware replanning module. Low risk executes the nominal trajectory; high risk (dashed branch) triggers the proposed replanning to ensure safety.

Optimization-based approaches, such as Model Predictive Control (MPC) and Dynamic Programming (DP), formulate driving decisions as constrained optimization problems. These methods minimize a cost function over a planning horizon while considering objectives such as safety, comfort, and traffic efficiency. MPC is particularly suitable for real-time decision-making, as it simultaneously optimizes vehicle control inputs (steering, acceleration, braking) and respects dynamic constraints including vehicle kinematics, road geometry, and the predicted behavior of other road users [4].

Prior work has explored a variety of methodologies for modeling occluded obstacles and integrating them into planning frameworks. For example, Batkovic et al. [5] present a safe MPC formulation incorporating pedestrian tracking and occlusion handling, building upon the Model Predictive Flexible Trajectory Tracking Control (MPFTC) framework [6]. Van der Ploeg et al. [7] employ risk fields derived from occupancy grids and reachability analysis to represent occluded objects continuously, enabling their integration as constraints in MPC. Fan et al. [8] propose an efficient collision-avoidance method using geometrical approaches within MPC, addressing both obstacles and constrained region boundaries. Park et al. [9] utilize a simplified reachability quantification to assess the risk associated with hidden pedestrians or vehicles, which is then incorporated into MPC for trajectory planning.

Collectively, these studies demonstrate that explicitly accounting for occlusions—whether via reachability analysis, risk fields, or geometrical constraints—is

critical for designing safe and reliable trajectory planners for autonomous vehicles.

1.2 Key Contributions

The main contributions of this thesis are as follows:

1. **Robust Trajectory Planning under Occlusions:** We develop a Robust Model Predictive Control (RMPC) framework that formally integrates worst-case, reachability-based predictions of hidden pedestrians and vehicles, enabling proactive collision avoidance in occluded urban environments.
2. **Comprehensive Simulation and Validation:** The proposed RMPC planner is rigorously evaluated in challenging urban scenarios, demonstrating a zero-collision rate and significant improvement over baseline methods.
3. **Safety–Performance Trade-off Analysis:** We systematically analyze the impact of robustness on key performance metrics such as traversal time and passenger comfort, providing insights into the trade-offs inherent in safety-aware planning.

1.3 Thesis Outline

The remainder of this thesis is organized as follows.

- **Chapter 2** provides the theoretical background, covering MPC, Robust MPC, collision avoidance via Hyperplane methods, vehicle modeling, reachability analysis, and the Frenet coordinate system.
- **Chapter 3** details the proposed methodology for modeling and predicting hidden agents, the computation of reachable sets, and the integration into the RMPC framework, including the objective function, constraints, and evaluation metrics.
- **Chapter 4** presents simulation results comparing the proposed occlusion-aware RMPC to a baseline approach in scenarios with hidden pedestrians and vehicles.
- **Chapter 5** concludes the thesis by summarizing key findings, discussing limitations, and suggesting directions for future work.

2

Background

This chapter lays the theoretical foundation for the proposed methodology by introducing the core principles and concepts. We begin with a general overview of Model Predictive Control (MPC), elucidating its fundamental structure and advantages as a robust control paradigm. Subsequently, we detail the extension to Robust Model Predictive Control (RMPC), which forms the bedrock of our approach, designed to handle uncertainties and disturbances inherent in real-world driving scenarios.

A crucial component of our collision avoidance strategy is the Hyperplane method, which we utilize to formulate explicit separation constraints within the MPC framework. The vehicle's dynamics are accurately represented by the Bicycle model, a widely accepted simplification that captures essential vehicular motion for efficient real-time control.

Furthermore, we elaborate on reachability analysis as a powerful tool for proactively modeling and predicting the potential positions of hidden road users, specifically addressing the uncertainties associated with occluded regions. Finally, the Frenet coordinate system is introduced as a pivotal element for trajectory representation, enabling more intuitive and efficient planning and control by decoupling longitudinal and lateral motions relative to a reference path.

2.1 Model predictive control

MPC is an advanced control method that uses a system model to optimize control over a time horizon while subject to different constraints and predict the future behavior of the system. MPC has been extensively used for trajectory planning as it has the ability to predict the future state of the system taking into consideration the model of the system while simultaneously satisfying different constraints. This is particularly important in the context of AD/ADAS where a vehicle is driving in a dynamic environment.

The further subsections discuss the formulation of an optimization problem as well as the formulation of Robust MPC that is made use of.

2.1.1 Basic Formulation of the optimization problem

The formulation involves an objective function J subjected to different constraints. The objective of the optimization problem is to either maximize or minimize the cost and is decided based on the application. There are different constraints such as equality, inequality constraints etc. depending on the formulation. Equality constraints are strict constraints that have to be satisfied while inequality constraints are more relaxed meaning there exists a boundary within which it has to be satisfied.

The equation (2.1) shows one such optimization problem where the goal is to minimize the cost. The typical constraints that are used in these problem formulation are shown in equation (2.2), (2.3), (2.4).

$$\min_{x} J(x) \tag{2.1}$$
$$s.t.\ h(x) = 0 \tag{2.2}$$
$$g(x) \leq 0 \tag{2.3}$$
$$x(t) \in \mathbb{X} \tag{2.4}$$

where

- J is the objective function
- x is the state
- h represents the equality constraints
- g represents the inequality constraints
- $\mathbb{X}$ represents the set of decision variables

2.1.2 Robust Model predictive control

Robust Model predictive control is a variant of MPC that explicitly accounts for disturbances or uncertainties described using w in equations (2.8). It ensures that the control actions remain feasible when it is subject to the bounded uncertainties.

$$\min_{x,u} V_f(x(t+T|t)) + \int_0^{t_f} V_c(x(t+\tau|t))dt \tag{2.5}$$
$$s.t.\ \dot{x}(t+\tau|t) = f(x(t+\tau|t), u(t+\tau|t)) \tag{2.6}$$
$$x(0) = x_0 \tag{2.7}$$
$$\forall w\ \in [\underline{w}, \overline{w}] \tag{2.8}$$
$$g(x, u, w) \leq 0 \tag{2.9}$$

where

- V_c is the running cost at time t
- V_f is the final cost at terminal time T

- x_0 is the initial state
- w is the disturbances
- t_f is the prediction time horizon

The equation involves two costs; one a running cost that adds up till the terminal time and the other a final cost at the terminal time.

2.2 Runge-Kutta Method

Runge-Kutta method is a popular technique that is used to solve Ordinary Differential Equations (ODEs). A commonly used version of this is the fourth order Runge-Kutta (RK4) and the equations are shown in (2.10).

$$\begin{aligned} k_1 &= f(x_n, u_n) \\ k_2 &= f\left(x_n + \frac{h}{2}, u_n + \frac{h}{2}k_1\right) \\ k_3 &= f\left(x_n + \frac{h}{2}, u_n + \frac{h}{2}k_2\right) \\ k_4 &= f(x_n + h, u_n + hk_3) \\ x_{n+1} &= u_n + \frac{h}{6}(k_1 + 2k_2 + 2k_3 + k_4) \end{aligned} \tag{2.10}$$

In the context of MPC, RK4 helps discretize the continuous time model. This method not works well for the equations but also provides high accuracy.

2.3 Hyperplane Separation Theorem

The hyperplane separation theorem states that, given two disjoint convex sets in Euclidean space there exists a hyperplane that separates them. In this problem statement the hyperplanes, shown in Figure. 2.1 are used to separate one object from the other object. Using this a safe boundary can be created that ensures that the ego vehicle doesn't collide with the other obstacles or road boundaries as demonstrated in Fan et al. in [8].

The equations (2.11), (2.12) describe the the two half spaces formed as a the result of the hyperplane splitting them.

$$H^+ = s \in \mathbb{R}^N | \lambda^T s \leq \mu \tag{2.11}$$

$$H^- = s \in \mathbb{R}^N | \lambda^T s \geq \mu \tag{2.12}$$

$$\lambda \in \mathbb{R}^N \tag{2.13}$$

Figure 2.1: Illustration of a hyperplane separating two polygons into two separate regions

Here H^+ and H^- are the regions that are formed as the result of hyperplane, λ is the normal vector to the hyperplane and μ is the intercept.

Using the above formulation the hyperplane equations used as constraints as described in the equation (2.14).

$$\lambda V_{P_1} \geq \mu, \lambda V_{P_2} \leq \mu, \lambda V_{Occ} \leq \mu \tag{2.14}$$

Here V_{P_1} is the polygon containing the ego vehicle, V_{P_2} is the polygon containing the static obstacle and V_{Occ} is the polygon representing the occluded obstacle.

2.4 Vehicle Model

As discussed in section 2.1, MPC requires a model of the system. There are different ways to choose a model depending on the application each offering certain advantages and disadvantages. The ego vehicle in this case is modeled using a kinematic bicycle model, shown in the Figure. 2.2. The kinematic bicycle model is a simple yet effective model that captures the necessary aspects of the vehicle. The rear axle (green point) is chosen to be the reference point for the formulation. Using the rear axle as the reference point not only simplifies calculations but also has stable dynamics.

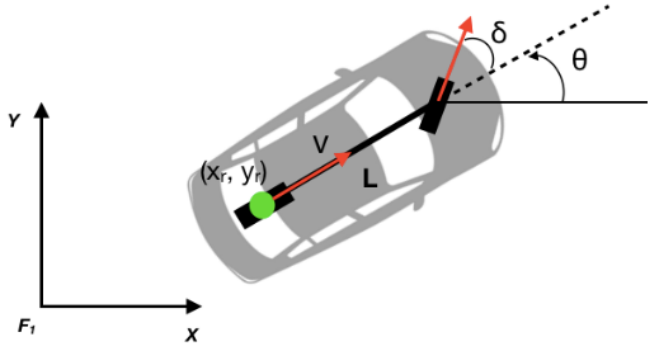

Figure 2.2: Vehicle Model

The state x and u is given in the equation (2.15)

$$\mathbf{x} = \begin{bmatrix} s \\ y \\ \theta \\ v \\ \delta \end{bmatrix}, \mathbf{u} = \begin{bmatrix} a \\ \omega \end{bmatrix} \tag{2.15}$$

The vehicle dynamics denoted by $\dot{x} = f(x, u)$ is modeled as shown in the equation (2.16)

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{s} \\ \dot{y} \\ \dot{\theta} \\ \dot{v} \\ \dot{\delta} \end{bmatrix} = \begin{bmatrix} v\cos\theta \\ v\sin\theta \\ \frac{v}{L}\tan\delta \\ a \\ \omega \end{bmatrix} \tag{2.16}$$

where

- (s, y) is the position of the center of rear axle of the vehicle.
- θ is the heading of the vehicle.
- δ is the steering angle.
- v is the velocity.
- a is the acceleration.
- ω is the steering rate.
- L is the wheelbase of the vehicle.

2.5 Reachability Analysis

Two primary concepts from reachability analysis: the Forward Reachable Set (FRS) for prediction and the Backward Reachable Set (BRS) for inference are leveraged here.

2.5.1 Forward Reachable Set

The Forward Reachable Set (FRS) encompasses all possible future states a system can reach from a given set of initial conditions over a specific time horizon, subject to its dynamics and control limitations. In the context of autonomous driving, the FRS is used to predict the future space that will be occupied by a agent.

Definition: Consider a dynamical system described by the ordinary differential equation (ODE) $\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t))$, where $\mathbf{x} \in \mathbb{R}^n$ is the state vector and $\mathbf{u} \in \mathbb{U}$ is the control input from a set of admissible controls $\mathbb{U}$. Given a set of initial states

$\mathbb{X}_0$ at time t_0, the FRS over the time interval $[t_0, t_f]$ is the set of all states that are reachable from any state in $\mathbb{X}_0$.

Mathematical Formulation: The FRS at a future time t_f, denoted as $\mathcal{R}_F(t_f; \mathbb{X}_0)$, is formally defined as:

$$\begin{aligned}\mathcal{R}_F(t_f; \mathbb{X}_0) = \{\mathbf{x}(t_f) \in \mathbb{R}^n \mid \exists \mathbf{x}_0 \in \mathbb{X}_0,\ \exists \mathbf{u}(\cdot) \in \mathbb{U}, \\ \text{s.t. } \dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t)) \text{ for } t \in [t_0, t_f]\}\end{aligned} \tag{2.17}$$

2.5.2 Backward Reachable Set

The Backward Reachable Set (BRS) is conceptually the inverse of the FRS. It identifies all initial states from which a system could have evolved to enter a specified target set $\mathbb{T}$ at a future time. Its primary power in our work is its ability to reason backwards from an area of concern—specifically, an occlusion.

Definition: For the same dynamical system, the BRS identifies all states $\mathbf{x}_0$ at time t_0 from which there exists at least one control trajectory $\mathbf{u}(t)$ that drives the system into a given target set $\mathbb{T}$ at or before a final time t_f.

Mathematical Formulation: The BRS, computed backwards from time t_f to t_0 for a target set $\mathbb{T}$, is defined as:

$$\begin{aligned}\mathcal{R}_B(t_0; \mathbb{T}) = \{\mathbf{x}(t_0) \in \mathbb{R}^n \mid\ \mathbf{x}(t_f) \in \mathbb{T}, \forall \mathbf{u}(\cdot) \in \mathbb{U}, \\ \text{s.t. } \dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t)) \text{ for } t \in [t_0, t_f]\}\end{aligned} \tag{2.18}$$

2.6 Frenet Coordinate System

In autonomous driving systems, trajectory planning and control present significant computational and theoretical challenges. While global Cartesian coordinates (X–Y) provide an intuitive geometric representation, they introduce substantial complexity in path-following applications due to the inherent coupling between longitudinal and lateral vehicle dynamics, particularly in curved road geometries. To address these fundamental limitations, the Frenet coordinate system [10] establishes a path-centric coordinate transformation that decouples vehicle motion relative to a prescribed reference trajectory, thereby facilitating more tractable control formulations.

2.6.1 Mathematical Foundation and Coordinate Transformation

The Frenet coordinate system provides a curvilinear description of a vehicle's state relative to a reference path, offering a natural parameterization for path-following and trajectory-planning problems. Consider a point (X, Y) in the global Cartesian frame. Its representation in the Frenet frame is uniquely specified by the triplet (s, d, θ_e):

- s is the arc-length coordinate, representing the cumulative distance along the reference path from a chosen origin to the orthogonal projection of (X, Y) onto the path.
- d is the signed lateral offset, defined as the perpendicular distance from the reference path to (X, Y), measured along the local normal vector at the projection.
- θ_e is the heading error, the angular difference between the vehicle heading θ and the tangent direction of the reference path θ_r at the projection,

$$\theta_e = \theta - \theta_r.$$

Let the reference path be parameterized by arc length s with position vector $(X_r(s), Y_r(s))$, tangential heading $\theta_r(s)$, and curvature $\kappa_r(s)$. To obtain the Frenet coordinates of a Cartesian state (X, Y, θ), determine the value of s that minimizes the Euclidean distance between (X, Y) and the reference path. With $(X_r(s), Y_r(s), \theta_r(s))$ denoting the path properties at this projection, the forward transformation is

$$\begin{aligned} d &= -(X - X_r(s)) \sin\Big(\theta_r(s)\Big) + (Y - Y_r(s)) \cos\Big(\theta_r(s)\Big), \\ \theta_e &= \theta - \theta_r(s). \end{aligned} \tag{2.19}$$

Conversely, the inverse transformation from Frenet coordinates (s, d, θ_e) back to Cartesian coordinates (X, Y, θ) is

$$\begin{aligned} X &= X_r(s) - d \sin\Big(\theta_r(s)\Big), \\ Y &= Y_r(s) + d \cos\Big(\theta_r(s)\Big), \\ \theta &= \theta_r(s) + \theta_e. \end{aligned} \tag{2.20}$$

These formulas follow directly from the geometry of the reference path and are commonly used in motion-planning and control for autonomous vehicles.

3

Methods

This chapter delves into the robust modeling and motion prediction of hidden agents using reachable set analysis. We specifically address two critical scenarios: hidden pedestrians and hidden vehicles, acknowledging their distinct dynamic characteristics and uncertainties, which necessitate tailored modeling approaches.

Our methodology encompasses the entire pipeline, beginning with the initial set inference for these hidden entities, proceeding through the detailed computation of their FRS. These rigorously determined reachable sets are then seamlessly integrated into a RMPC framework.

Beyond this integration, the chapter provides a comprehensive exposition of the proposed RMPC's specific design and core components. This includes a detailed examination of the objective function, meticulously designed to balance the often-conflicting objectives of performance, comfort, and safety, alongside the constraints that ensure physical feasibility, operational limits, and adherence to traffic regulations. Finally, a precise overview of the simulation parameters employed to validate the efficacy of our approach is presented.

Ultimately, to provide a holistic assessment of the generated trajectories, three distinct metrics are introduced, evaluating performance across the critical dimensions of safety, efficiency, and comfort.

3.1 Modeling of Road Users

There are two methods used here namely a Simulated Road User and a Reachability set. The simulated road user is a point mass model that has a lateral and longitudinal velocity and the reachability set is calculated using the minimum and maximum expected velocity of the road user.

3.1.1 Simulated Road User

This model assumes the road user as a point mass with different velocities both laterally and longitudinally. The longitudinal velocity values of the road user in this case a pedestrian were chosen in the range of 0 to 1.5 m/s and between -0.5 and 0.5 m/s laterally.

Using the range described before, the velocities were chosen using two different methods

- **Uniform distribution** The values were picked randomly from a uniform distribution in the range of the expected velocity for a certain road user as shown in figure 3.1. This means that each value has an equal possibility of being chosen.

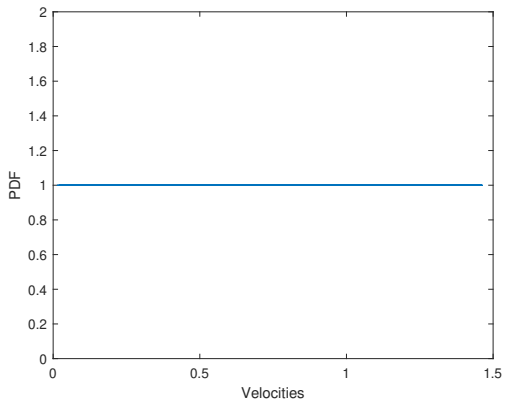

Figure 3.1: Uniform used to pick velocities

- **Skewed Probabilistic distribution** The values were picked randomly from a beta distribution shown in figure 3.2 in the range of the expected velocity for a certain road user. By using a beta distribution the velocity values can instead be chosen from a selective area within the given range. The values also have different possibilities of occurring as shown in the figure 3.2 A beta distribution typically has two controlled parameters alpha (α and β). The effect of these parameters are described below.

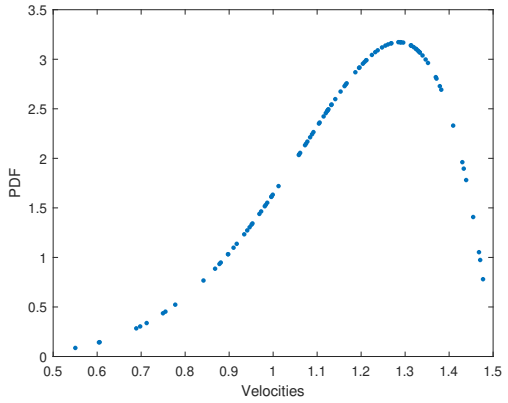

Figure 3.2: Beta Distribution used to pick velocities

  - When $\alpha > \beta$ the graph is skewed to the right meaning more values close to upper bound

- $\alpha < \beta$ the graph is skewed to the left meaning more values close to lower bound
- $\alpha = \beta = 1$ gives a uniform distribution

3.2 Motion Prediction and Inference for Hidden Pedestrians and Vehicles

Trajectory planning under occlusion presents a fundamental challenge in autonomous driving: managing uncertainty about unobserved agents. The occluded regions introduce critical unknowns—whether pedestrians or vehicles are present, their current states, and their intended motions. This uncertainty necessitates a principled approach that can reason about all possible threat scenarios while maintaining computational tractability.

Reachability analysis [11] provides a rigorous mathematical framework for computing the set of all possible states that a dynamical system can occupy over time. This methodology is fundamental to our approach, enabling both the prediction of visible agent behaviors and, more critically, the systematic inference of potential threats from hidden agents in occluded regions. Unlike heuristic methods that rely on simplified assumptions, reachability analysis offers formal guarantees about safety and completeness of the threat assessment.

By integrating both FRS and BRS, we construct a comprehensive risk representation that accounts for the full spectrum of possible hidden agent behaviors. This integration enables robust trajectory planning that maintains safety even under complete uncertainty about occluded regions.

Our methodology consists of three interconnected components:

1. **Initial Set Inference for Hidden Agents:** For each occluded scenario, we first determine the feasible initial state set $\mathbb{P}_0$ from which hidden agents could potentially emerge to threaten the ego vehicle. The inference methodology differs fundamentally between pedestrians and vehicles: hidden pedestrians may remain stationary within occlusions and exhibit diverse motion patterns, while hidden vehicles follow structured lane-based dynamics with predictable motion constraints.

2. **Forward Reachable Set Computation:** Given the inferred initial sets, we compute the time-dependent FRS that characterizes all possible future states of hidden agents. This computation accounts for the dynamic interaction between the ego vehicle's motion and the evolving threat landscape as occlusions change.

3. **Safety-Constrained Trajectory Optimization:** The predicted FRS are incorporated as collision avoidance constraints within the RMPC framework. This integration ensures that the planned trajectory maintains safety margins

against all possible hidden agent behaviors while optimizing for efficiency and comfort.

The following subsections detail the specific methodologies for hidden pedestrian and hidden vehicle scenarios, highlighting the distinct challenges and solutions for each agent type. We demonstrate how the unified reachability framework adapts to the unique characteristics of different road users while maintaining formal safety guarantees.

3.2.1 Motion Prediction for Hidden Pedestrians

3.2.1.1 Hidden Pedestrian Scenario Formulation

In complex urban driving environments, static obstacles such as parked vehicles, roadside infrastructure, and buildings frequently obstruct sensor observations, creating blind spots where pedestrians may be concealed. The phenomenon of pedestrians suddenly emerging from occlusions—commonly referred to as "ghost pedestrian" scenarios—represents one of the most challenging problems in autonomous driving safety.

We consider a canonical hidden pedestrian scenario as illustrated in Figure. 3.3. The ego vehicle navigates along a roadway where three parked vehicles create significant occlusions adjacent to the travel lane. These occluded regions may conceal pedestrians who could potentially emerge and cross the roadway at any moment, creating imminent collision threats. The planning system must account for these invisible agents and generate trajectories that maintain safety margins against all possible pedestrian emergence scenarios.

The fundamental challenge lies in the complete absence of observational data about potential hidden pedestrians. Unlike visible agents whose states can be observed and tracked, hidden pedestrians introduce pure uncertainty: their presence, initial states, intended motions, and emergence timing are all unknown. This necessitates a worst-case analysis approach that considers all plausible threat scenarios while maintaining computational tractability for real-time implementation.

3.2.1.2 Initial State Set Inference for Hidden Pedestrians

The inference of plausible initial pedestrian states begins with identifying the target region where a hidden pedestrian could create a collision threat with the ego vehicle's planned trajectory. We define this target set $\mathbb{T}$ as the visible road segments that lie within the ego vehicle's future path and where pedestrian-vehicle conflicts could occur.

Given the target set $\mathbb{T}$, we employ backward reachable set analysis to determine all initial positions from which a pedestrian could potentially reach $\mathbb{T}$ within the relevant time horizon. The backward reachable set $\mathbb{B}$ captures all states from which the pedestrian dynamics can reach the target region, accounting for the pedestrian's motion capabilities and the available time window.

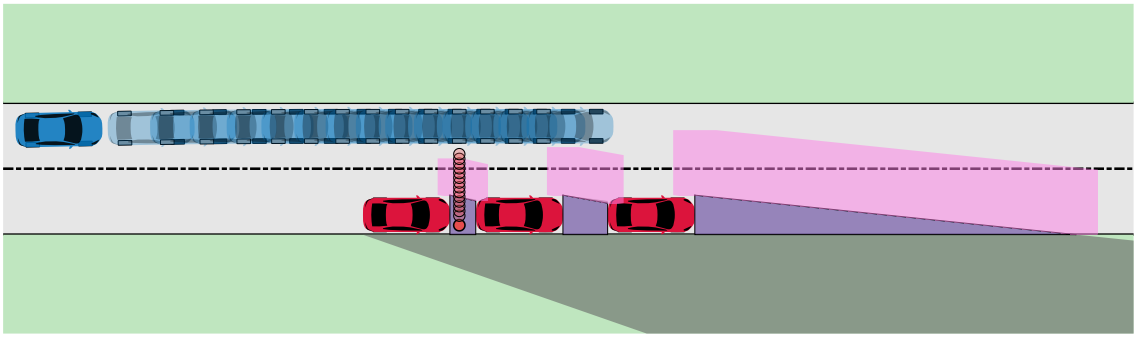

Figure 3.3: Hidden pedestrian scenario: ego vehicle (blue) traverses a road segment with three parked vehicles (red). Gray regions represent occlusions caused by parked vehicles, purple regions denote inferred initial sets of potential hidden pedestrians, pink areas show forward reachable sets, and transparent ego vehicles indicate predicted future positions.

The feasible initial set $\mathbb{X}_0$ for hidden pedestrians is then computed as the intersection of the backward reachable set with the occluded region:

$$\mathbb{X}_0 = \mathbb{B} \cap \mathbb{O}_{occ} \tag{3.1}$$

where $\mathbb{O}_{occ}$ represents the geometrically characterized occlusion region based on sensor field-of-view limitations and environmental obstacles. The resulting set $\mathbb{X}_0$ represents all plausible initial positions of hidden pedestrians who could potentially create conflict scenarios with the ego vehicle's nominal trajectory. Excluding the region outside of the road boundaries, this set is visualized as the purple convex polygon in Figure. 3.3.

This methodology ensures that the initial set captures all safety-critical scenarios while remaining computationally tractable. The approach is formalized in Algorithm 1.

Algorithm 1 Initial State Set Inference for Hidden Pedestrians

1: **procedure** InferPedestrianInitialSet
2: **Define Target Set** $\mathbb{T}$: Specify the visible region on the road ahead where pedestrian-vehicle conflicts could occur
3: **Compute BRS**: Calculate the Backward Reachable Set $\mathbb{B}$ of the pedestrian model with objective to reach $\mathbb{T}$
4: **Define Occlusion Geometry** $\mathbb{O}_{occ}$: Geometrically characterize the occluded area based on sensor limitations and environmental obstacles
5: **Compute Initial State Set** $\mathbb{X}_0$: Calculate the intersection between the BRS and the occluded region:

$$\mathbb{X}_0 = \mathbb{B} \cap \mathbb{O}_{occ} \tag{3.2}$$

6: **return** $\mathbb{P}_0$

3.2.1.3 Forward Reachable Set Computation for Hidden Pedestrians

Once the initial state set $\mathbb{X}_0$ is established, we must characterize the future evolution of potential hidden pedestrians under uncertainty. Since we have no observational data about hidden pedestrians—including their presence, current states, or intended motions—we adopt a robust worst-case approach that considers all plausible pedestrian behaviors within each RMPC prediction horizon.

We model the pedestrian as a point mass with bounded velocity constraints. The state vector is defined as $\mathbf{x}_{ped} = [x, y]^T$, representing the pedestrian's position in Cartesian coordinates. The pedestrian dynamics are governed by $\dot{\mathbf{x}}_{ped} = [v_x, v_y]^T$, where the velocity components are constrained within physically realistic bounds that capture the full range of pedestrian motion capabilities. To ensure comprehensive threat coverage, we model the pedestrian velocities as uniformly distributed over bounded intervals: $v_x \in [-v^{ped}_{x,\max}, v^{ped}_{x,\max}]$ and $v_y \in [0, v^{ped}_{y,\max}]$. The lateral velocity constraint v_x allows bidirectional movement parallel to the roadway, reflecting the pedestrian's ability to move both toward and away from the ego vehicle's path. The longitudinal velocity constraint v_y is unidirectional and positive, focusing on road-crossing behaviors that create collision threats. We set $v^{ped}_{x,\max} < v^{ped}_{y,\max}$ to reflect the higher likelihood of road-crossing versus lateral movement patterns.

The forward reachable set computation accounts for the time-varying nature of occlusions as the ego vehicle progresses. For each occluded region, the available propagation time t for hidden pedestrians is determined by $t = d_{occ}/v_{ego}$, where d_{occ} represents the distance from the ego vehicle's current position to the point where the occlusion becomes fully visible, and v_{ego} is the ego vehicle's velocity along its nominal trajectory. Figure 3.4 illustrates the forward reachable set propagation at fixed time intervals of 0.6 s across the RMPC prediction horizon, demonstrating how uncertainty grows over time and requires increasingly conservative safety margins for longer prediction horizons.

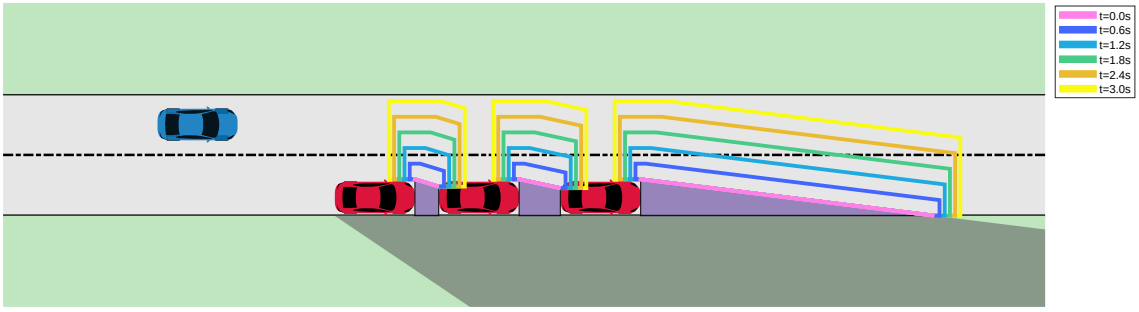

Figure 3.4: Forward reachable sets of hidden pedestrians propagated in 0.6 s time steps across the RMPC prediction horizon, demonstrating uncertainty growth over time.

Figure 3.5 demonstrates the evolution of forward reachable sets as the ego vehicle progresses through the occluded scenario. Several key observations emerge from this analysis. First, pedestrians hidden behind the furthest vehicles present larger potential reachable sets due to longer available propagation times, creating distance-

dependent uncertainty in threat estimation that demands more cautious consideration of potential pedestrian trajectories. Second, as the ego vehicle approaches the occluded regions, the available propagation time decreases, leading to smaller reachable sets and reduced collision probability, demonstrating progressive threat reduction. Finally, when the ego vehicle passes each parked vehicle, the corresponding occluded region becomes fully observable, eliminating the associated hidden pedestrian threat through occlusion clearance.

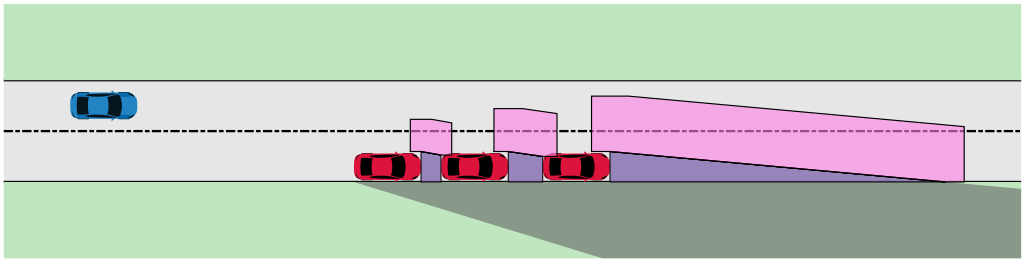

(a) $t = 2.4\,\mathrm{s}$: Maximum threat scenario with largest reachable sets

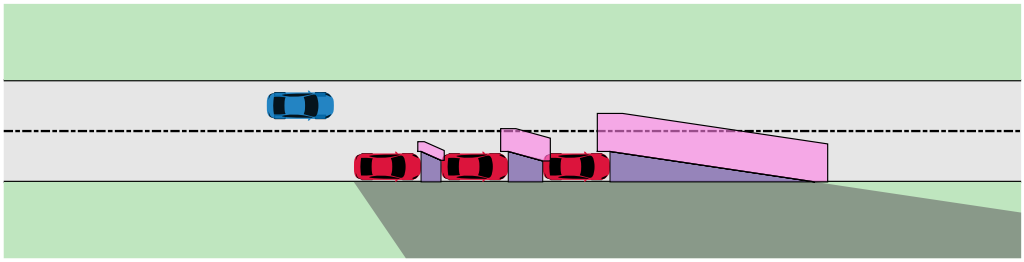

(b) $t = 3.6\,\mathrm{s}$: Intermediate stage with reduced reachable sets

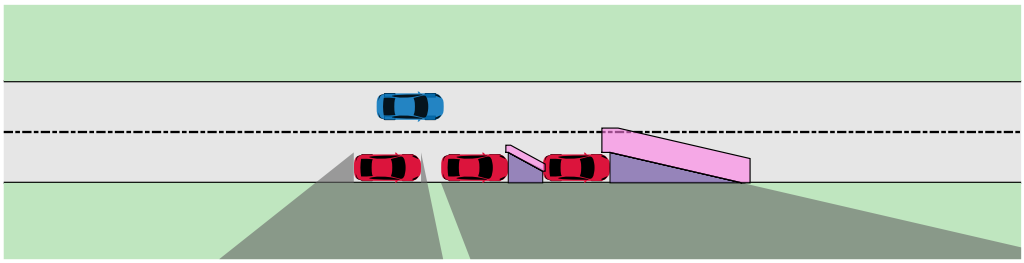

(c) $t = 4.4\,\mathrm{s}$: Near-complete occlusion clearance with minimal threats

Figure 3.5: Spatial-temporal evolution of hidden pedestrian forward reachable sets. The progressive shrinkage and elimination of reachable sets demonstrates threat reduction as the ego vehicle approaches and passes occluded regions.

The computed forward reachable set provides a formal over-approximation of all spatial regions that hidden pedestrians might occupy over the prediction horizon. This probabilistically complete threat characterization serves as the foundation for safety constraint generation in the RMPC framework, ensuring that planned trajectories maintain adequate safety margins against all considered pedestrian emergence scenarios while preserving system performance and computational efficiency.

3.2.2 Motion Prediction for Hidden Vehicles

3.2.2.1 Hidden Vehicle Scenario Formulation

Navigating complex urban intersections represents one of the most challenging scenarios in autonomous driving, primarily due to severe sensor occlusions induced by roadside infrastructure and environmental structures. Unlike open highway scenarios where visibility extends far into the horizon, urban intersections create multiple blind spots through buildings, traffic infrastructure, and other vehicles, fundamentally limiting the ego vehicle's perception capabilities.

As illustrated in Figure. 3.6, when an ego vehicle executes a left turn maneuver at an intersection, its perception of crossing traffic lanes is frequently obstructed by buildings and other static structures (gray rectangles). These obstructions create extensive occluded regions (gray areas) from which another vehicle could emerge unexpectedly, presenting significant and imminent collision risks. The critical challenge lies in the fact that vehicles in cross-traffic lanes may be approaching at substantial velocities while remaining completely invisible to the ego vehicle's sensor suite until they emerge from the occlusion.

To guarantee safety-critical operation under such uncertainty, the trajectory planning system must proactively account for the potential presence and motion of hidden vehicles. This necessitates a systematic approach that identifies the most critical occluded regions, infers plausible initial vehicle states, and generates conservative motion predictions for potential threats operating within these blind spots.

3.2.2.2 Risk-Aware Occlusion Analysis

A fundamental insight in occlusion-aware planning is that not all occluded regions pose equivalent levels of threat to the ego vehicle. Treating all occlusions equally would result in excessive conservatism and computational burden, potentially rendering the planning system impractical for real-time applications. To address this challenge, we develop a risk-aware classification framework that systematically identifies and prioritizes occluded areas based on their collision potential with the ego vehicle's planned trajectory.

Our approach builds upon visibility-based occlusion analysis methodologies [9], which examine the intersection of visibility polygons with road network topology. However, we propose a more direct classification framework based on the ego vehicle's motion capabilities and temporal constraints, enabling efficient computation of safety-critical scenarios.

Lane Intersection Node Identification: The first step involves determining the critical boundary points where the ego vehicle's line-of-sight is interrupted along the centerlines of relevant traffic lanes. Let $\mathbb{O}_{obs}$ denote the polygon representing the currently observable area (visibility polygon) from the ego vehicle's sensor configuration, and let $\{\mathcal{L}_k\}_{k=1}^{m}$ represent the centerlines of relevant traffic lanes. The set of all critical intersection nodes, $\mathbb{N}$, is formally defined as:

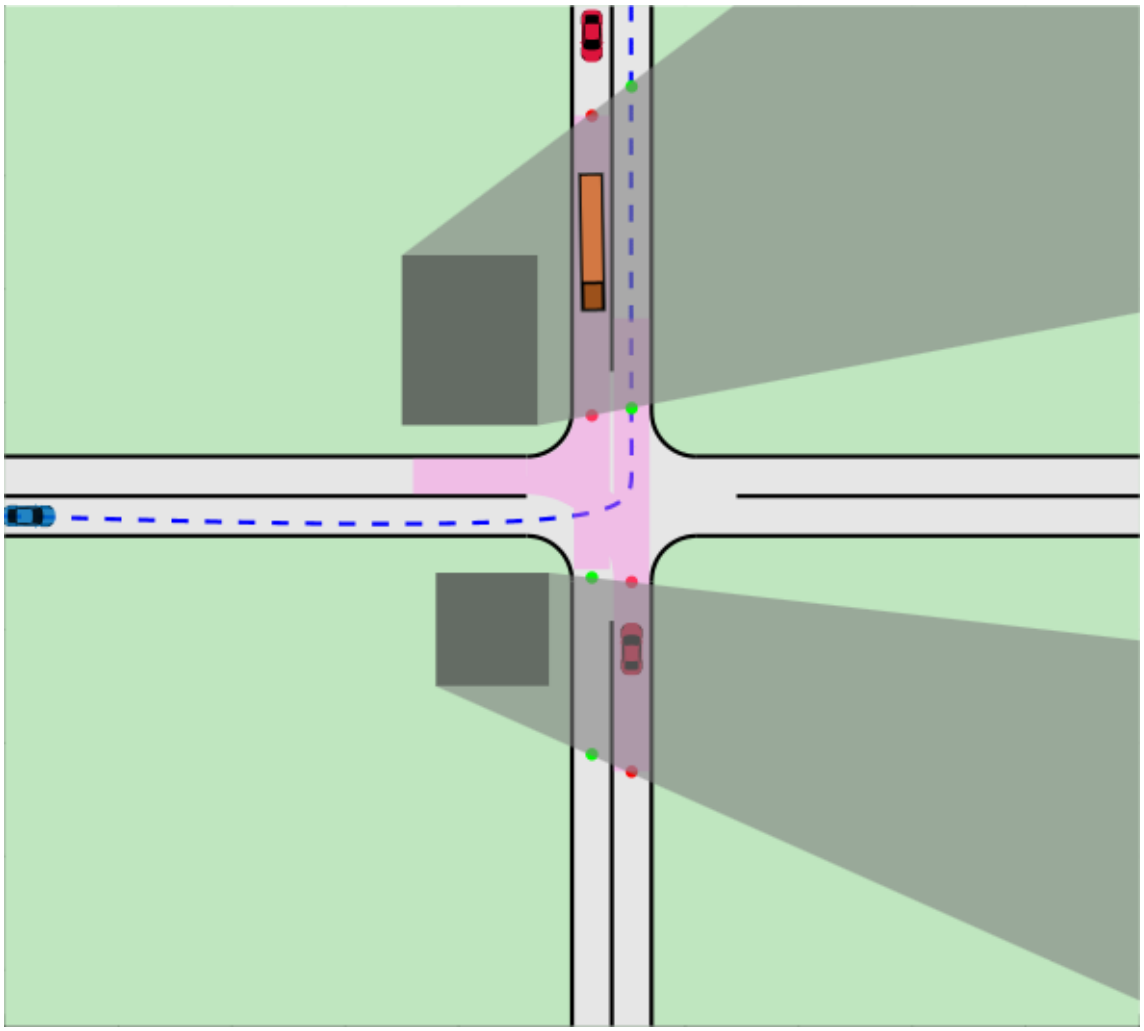

Figure 3.6: Hidden vehicle scenario: ego vehicle (blue) approaches a left turn at an intersection where buildings (gray rectangles) occlude portions of the crossing traffic lanes. Gray regions represent occluded areas, green solid circles indicate low-risk intersection nodes, red solid circles denote high-risk nodes, and pink regions represent potential positions of hidden vehicles in the future.

$$\mathbb{N} = \left\{ \mathbf{p} \in \mathbb{R}^2 \mid \mathbf{p} \in \partial\mathbb{O}_{obs} \cap \mathcal{L}_k, \forall k \in \{1, \ldots, m\} \right\} \tag{3.3}$$

where $\partial\mathbb{O}_{obs}$ denotes the boundary of the visibility polygon, and m represents the total number of relevant traffic lanes in the intersection vicinity.

Risk-Aware Node Classification: Subsequently, these intersection nodes are classified into low-risk and high-risk categories based on their potential for creating collision scenarios with the ego vehicle's planned trajectory. The fundamental principle underlying this classification is that a node only constitutes a high risk if the ego vehicle could physically reach a conflicting spatial location within the prediction horizon, thereby creating a potential collision scenario.

To formalize this concept, we compute the Forward Reachable Set (FRS) of the ego vehicle, $\mathcal{R}_F(t_f)$, over a prediction horizon t_f. A node $\mathbf{p} \in \mathbb{N}$ is classified as high-risk if a potential collision involving a vehicle emerging from that location could occur within the planning horizon. The high-risk node set is defined as:

$$\mathbb{N}_{high} = \{ \mathbf{p} \in \mathbb{N} \mid \text{isConflict}(\mathbf{p}, \mathcal{R}_F(t_f)) \} \tag{3.4}$$

where the function $\text{isConflict}(\cdot, \cdot)$ evaluates whether a vehicle emerging from node $\mathbf{p}$ could potentially lead to a collision with any state in $\mathcal{R}_F(t_f)$. The complementary low-risk node set is simply $\mathbb{N}_{low} = \mathbb{N} \setminus \mathbb{N}_{high}$.

As visualized in Figure. 3.6, high-risk nodes are marked in red, while low-risk nodes are indicated in green. For computational efficiency, the low-risk nodes and their associated occluded regions are excluded from further analysis. A critical observation is that occluded lane segments typically yield pairs of high-risk nodes, which can be classified as "near" and "far" nodes based on their distance to the ego vehicle.

3.2.2.3 Dynamic Initial Set Inference for Hidden Vehicles

Unlike hidden pedestrian scenarios, where pedestrians might remain stationary within occluded regions for extended periods before emerging suddenly, hidden vehicles present fundamentally different behavioral characteristics. It is highly improbable for a vehicle to remain parked motionless within an active intersection and then suddenly initiate movement. Instead, a vehicle's presence in cross-traffic lanes implies continuous motion governed by structured traffic dynamics.

This behavioral insight enables a more sophisticated and less conservative approach to initial set inference. Rather than maintaining static threat regions, we can dynamically update and potentially shrink the initial sets of hidden vehicles based on the evolving visibility conditions and the absence of observed vehicles.

Dynamic Node Evolution: The evolution of hidden vehicle threat regions is conceptualized through the movement of boundary nodes. For each occluded lane segment bounded by a pair of high-risk nodes, we define:

- **Near Node $\mathbf{p}^{\text{near}}(t)$**: Represents the closest possible position of a hidden vehicle to the ego vehicle at time t. This node is continuously updated based on current visibility-lane intersections.
- **Far Node $\mathbf{p}^{\text{far}}(t)$**: Represents the farthest possible initial position from which a hidden vehicle could have originated. This node propagates from its initial position at minimum velocity $v_{\min}^{veh}$.

The key distinction in initial conditions is crucial:

$$\mathbf{p}^{\text{far}}(0) = \mathbf{p}_0^{\text{far}} \quad \text{(absolute beginning of occluded area)} \tag{3.5}$$

$$\mathbf{p}^{\text{near}}(t) = \text{CurrentNearIntersection}(\partial\mathbb{O}_{obs}(t), \mathcal{L}_r) \quad \text{(current visibility boundary)} \tag{3.6}$$

Conservative Threat Elimination: A critical aspect of this approach is its ability to reduce conservatism systematically. When the arc length between the near and far nodes becomes insufficient to contain even the smallest vehicle, the occluded area can be definitively considered clear of hidden vehicles:

$$\text{ArcLen}(\mathbf{p}^{\text{near}}(t), \mathbf{p}^{\text{far}}(t)) < L_{veh} + d_{safe} \Rightarrow \mathcal{V}_0 = \emptyset \tag{3.7}$$

This calculation ensures that we maintain appropriate safety margins without excessive conservatism, allowing the system to confirm occlusion safety even before complete visibility restoration.

Figure. 3.7 illustrates the dynamic evolution of initial sets across different time instances, demonstrating the progressive refinement of threat assessment:

1. **Initial Assessment** (Figure. 3.7a, $t = 0\,\text{s}$): Red filled and hollow circles represent the initial near and far nodes, respectively, with pink rectangles indicating the inferred initial sets.
2. **Progressive Refinement** (Figure. 3.7b, $t = 2.5\,\text{s}$): As the ego vehicle advances, the initial sets shrink due to improved visibility and node evolution.
3. **Partial Clearance** (Figure. 3.7c, $t = 2.9\,\text{s}$): The lower occluded region's threat is eliminated, even though significant occlusions remain visible. This demonstrates the algorithm's ability to confirm safety without complete occlusion clearance.
4. **Near-Complete Clearance** (Figure. 3.7d, $t = 4.0\,\text{s}$): The upper road's threat is also eliminated, confirming the intersection's safety well before complete visibility restoration.

The complete algorithm for dynamic initial set inference is presented in Algorithm 2.

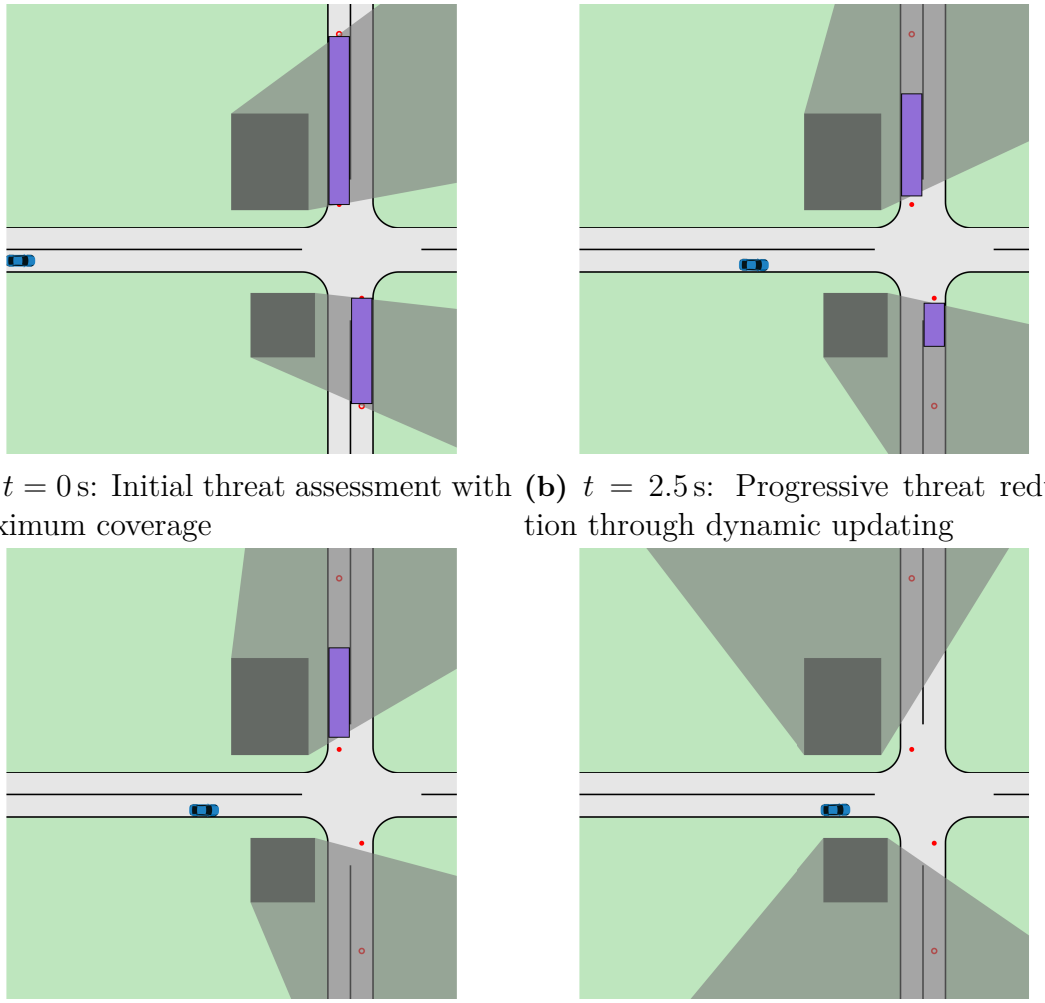

(a) $t = 0\,\mathrm{s}$: Initial threat assessment with maximum coverage

(b) $t = 2.5\,\mathrm{s}$: Progressive threat reduction through dynamic updating

(c) $t = 2.9\,\mathrm{s}$: Lower corridor confirmed safe

(d) $t = 4.0\,\mathrm{s}$: Upper corridor confirmed safe

Figure 3.7: Dynamic evolution of hidden vehicle initial sets. Red filled and hollow circles represent near and far nodes at $t = 0$, respectively. Blue rectangles denote the inferred initial sets, which vanish through systematic threat elimination before complete occlusion clearance.

Algorithm 2 Dynamic Initial Set Inference for Hidden Vehicles

Require: Visibility polygon $\mathbb{O}_{obs}$; lane centerlines $\{\mathcal{L}_k\}_{k=1}^{m}$; ego FRS $\mathcal{R}_F(t_f)$; vehicle envelope (L_{veh}, W_{veh}); lane width W_{lane}; minimum speed $v_{\min}$; time step Δt; safety margin d_{safe}.

Ensure: Inferred initial hidden vehicle set $\mathbb{I}_{veh}$.

Lane Intersection Node Identification:

$\mathbb{N} \leftarrow \{\, \mathbf{p} \in \mathbb{R}^2 \mid \mathbf{p} \in \partial\mathbb{O}_{obs} \cap \mathcal{L}_k,\ k = 1, \ldots, m \,\}$

Risk-Aware Node Classification:

$\mathbb{N}_{high} \leftarrow \{\, \mathbf{p} \in \mathbb{N} \mid \text{isConflict}(\mathbf{p}, \mathcal{R}_F(t_f)) \,\}$

Node Grouping by Occluded Corridors:

Partition $\mathbb{N}_{high}$ into groups $\{\mathcal{G}_r\}$, each corresponding to one occluded lane corridor.

Initialization:

$\mathbb{I}_{veh} \leftarrow \emptyset$

for all $\mathcal{G}_r$ **do**

  Order $\mathcal{G}_r$ along lane progression (increasing arc-length s).

  Form consecutive node pairs $(\mathbf{p}_0^{\text{near}}, \mathbf{p}_0^{\text{far}})$ within $\mathcal{G}_r$.

  for all pair $(\mathbf{p}_0^{\text{near}}, \mathbf{p}_0^{\text{far}})$ **do**

    $t \leftarrow 0$; $\mathbf{p}^{\text{far}}(0) \leftarrow \mathbf{p}_0^{\text{far}}$; $\mathcal{C} \leftarrow \emptyset$

    while $\text{ArcLen}(\mathbf{p}^{\text{near}}(t), \mathbf{p}^{\text{far}}(t)) \geq L_{veh} + d_{safe}$ **do**

      ▷ (1) Update near node from current visibility-lane intersection

      $\mathbf{p}^{\text{near}}(t) \leftarrow \text{CurrentNearIntersection}(\partial\mathbb{O}_{obs}, \mathcal{L}_r)$

      ▷ (2) Advance far node from previous position at $v_{\min}$

      $\mathbf{p}^{\text{far}}(t{+}\Delta t) \leftarrow \text{AdvanceToward}\big(\mathbf{p}^{\text{far}}(t),\ \mathbf{p}^{\text{near}}(t),\ v_{\min}^{veh}\Delta t\big)$

      ▷ Prevent far node from crossing near node

      $\mathbf{p}^{\text{far}}(t{+}\Delta t) \leftarrow \text{ProjectUpToNear}(\mathbf{p}^{\text{far}}(t{+}\Delta t),\ \mathbf{p}^{\text{near}}(t))$

      ▷ (3) Accumulate hidden corridor slice at time t

      $\mathcal{C} \leftarrow \mathcal{C} \cup \text{RectAroundCenterline}\big(\mathbf{p}^{\text{near}}(t),\ \mathbf{p}^{\text{far}}(t),\ \frac{1}{2}W_{lane}\big)$

      $t \leftarrow t{+}\Delta t$

    $\mathbb{I}_{veh} \leftarrow \mathbb{I}_{veh} \cup \mathcal{C}$

return $\mathbb{I}_{veh}$

3.2.2.4 Forward Reachable Set Computation for Hidden Vehicles

Once the dynamic initial sets are established, the computation of FRS for hidden vehicles requires careful consideration of vehicle dynamics and traffic constraints. Unlike pedestrians, which can exhibit arbitrary motion patterns, vehicles should typically follow structured lane-based behaviors.

For intersection scenarios with potential conflicts along straight road segments, we model hidden vehicles using constant-velocity dynamics along lane centerlines departing from near nodes. The hidden vehicle state is represented as $\mathbf{x}_{\text{veh}} = [s, v]^T$, where s denotes the arc-length position along the lane centerline and v represents the constant longitudinal velocity. Under the constant-velocity assumption, the system dynamics reduce to $\dot{\mathbf{x}}_{\text{veh}} = [v, 0]^T$.

To ensure comprehensive threat coverage, hidden vehicle speeds are modeled as uniformly distributed over the interval $[0, v_{\max}^{\text{veh}}]$, where $v_{\max}^{\text{veh}}$ represents the road speed limit. This uniform distribution captures the full spectrum of feasible driving speeds in urban intersection environments. Vehicles are assumed to occupy 90% of the lane width to account for realistic lateral positioning variations. While cross-traffic vehicles may choose to proceed straight or turn at the intersection, we focus on straight-line motion prediction for computational efficiency and worst-case threat representation.

The computed forward reachable sets provide a formal characterization of all potential spatio-temporal regions that hidden vehicles might occupy. As illustrated in Figure 3.8, the pink boundary represents the initial position from which the forward reachable set computation begins. The figure demonstrates the progressive expansion of reachable sets at 0.6-second intervals, with each temporal layer capturing the evolving threat envelope of potential hidden vehicle locations. Upon reaching the prediction horizon, the yellow boundary delineates the maximum extent of the forward reachable set, encompassing all feasible vehicle positions within the considered time window. This comprehensive spatio-temporal threat representation enables the RMPC framework to generate trajectories that maintain safety guarantees against all considered hidden vehicle scenarios while simultaneously optimizing for navigation efficiency and passenger comfort.

3.3 Setup of Occlusion-aware RMPC

This section details the occlusion-aware RMPC formulation employed for our proposed trajectory planning method. We delineate the chosen vehicle dynamics model, the design of the objective function, the definition of various constraints, and the key simulation parameters.

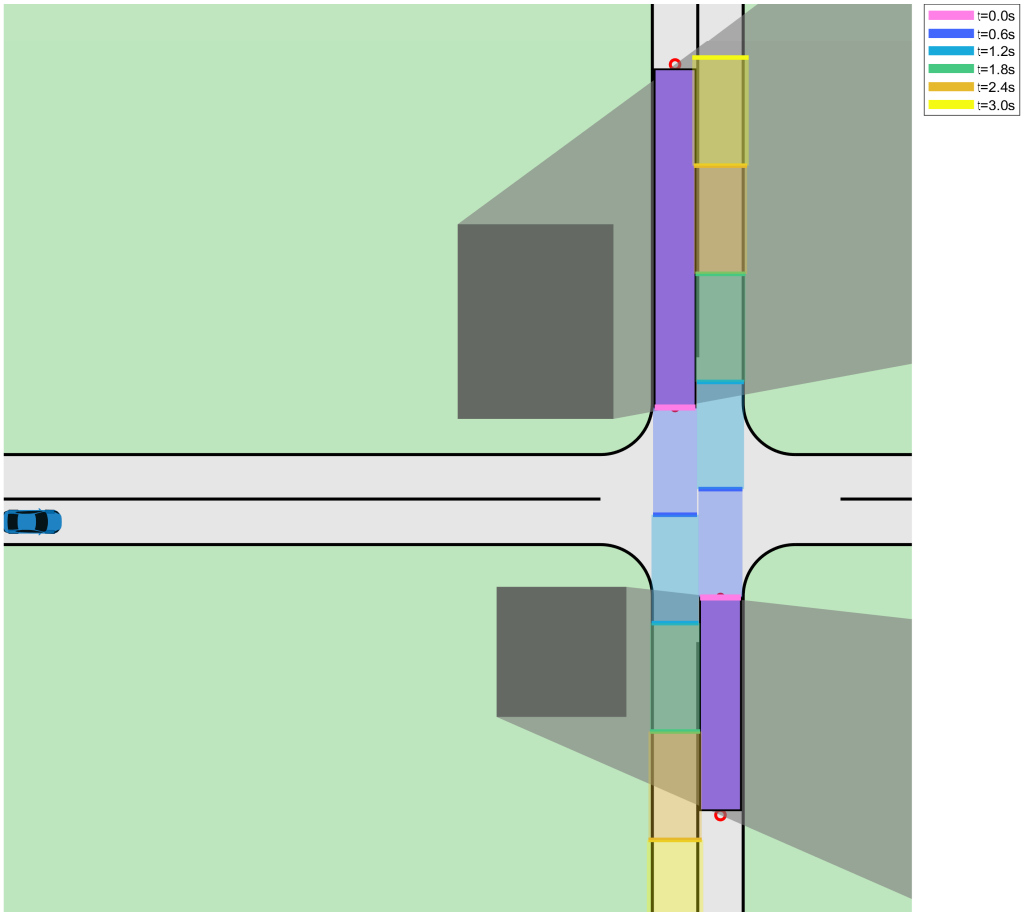

Figure 3.8: FRS expansion for hidden vehicles. Hidden vehicles originate from the boundary of near nodes within the initial set (purple region). The pink boundary represents the FRS at $t = 0$. The figure illustrates the progressive expansion of reachable sets at 0.6-second intervals, with each temporal layer depicted in different colors. The yellow region corresponds to the FRS expansion during the final 0.6 seconds of the prediction horizon. The complete FRS over the entire prediction horizon spans from the pink boundary to the yellow boundary.

3.3.1 Objective Function Design

The total objective function J of the proposed RMPC is formulated as:

$$J = \sum_{n=1}^{N-1} \ell_x(x_{k+n}, r_{k+n}) + \sum_{n=0}^{N-1} \ell_u(u_{k+n}) + \ell_{\text{terminal}}(x_{k+N}, r_{k+N}) + \ell_{\text{slack}} \tag{3.8}$$

where x_k and u_k represent the state and control input at prediction step k, $r_{\text{ref},k}$ is the reference state, $\ell(\cdot)$ quadratically penalizes the items accordingly.

3.3.1.1 Stage Cost

The stage cost, including state tracking costℓ_x and control input effort costℓ_u, penalizes deviations from the reference trajectory and excessive control effort respectively at each time step. The state tracking cost is defined as:

$$\ell_x(x_k, r_k) = (x_k - r_k)^T Q (x_k - r_k) \tag{3.9}$$

The diagonal elements of matrix Q are:

$$Q = \text{diag}([w_s, w_d, w_{\theta_e}, w_v, w_\delta, w_a]) \tag{3.10}$$

where w_s is weight for arc length progress, w_d is weight for lateral deviation, w_{θ_e} is weight for heading error, w_v is weight for longitudinal velocity tracking, w_δ is weight for steering angle and w_a is weight for longitudinal acceleration.

The control input cost can be defined as:

$$\ell_u(u_k) = u_k^T R u_k \tag{3.11}$$

The diagonal elements of R matrix are:

$$R = \text{diag}([w_{\text{jerk}}, w_\omega]) \tag{3.12}$$

where w_{jerk} is weight for longitudinal jerk and w_ω is weight for steering rate.

3.3.1.2 Terminal Cost

The terminal cost $\ell_{\text{terminal}}(x_N, x_{\text{ref},N})$ is applied at the end of the prediction horizon ($k = N$). It encourages the vehicle's state to converge towards the reference state, contributing to the stability of the MPC. It typically takes a similar quadratic form as the state tracking part of the stage cost, possibly with higher weights.

$$\ell_{\text{terminal}}(x_N, r_{\text{k+N}}) = (x_{k+N} - r_{\text{k+N}})^\top Q_{\text{term}} (x_{k+N} - r_{\text{x+N}}) \tag{3.13}$$

3.3.1.3 Slack Cost

Slack variables (S) are introduced for critical constraints, particularly for collision avoidance. While these constraints are primarily enforced as hard constraints, the use of slack variables allows for a degree of violation in situations where a hard constraint would render the problem infeasible. A large quadratic penalty is imposed on these slack variables to minimize their activation:

$$\ell_{\text{slack}} = \sum_{k=0}^{N-1} S_k^\top R_S S_k \tag{3.14}$$

where $R_S = \text{diag}([w_{S_1}, w_{S_2}, \ldots, w_{S_n}])$ defines the penalty weights for each slack variable, with particularly high weights for avoidance of hidden agents' FRS to ensure safety.

3.3.2 Constraints Design

Constraints are crucial for ensuring MPC operates within physical limits, adheres to traffic rules, and avoids collisions. These are categorized into hard and soft constraints.

3.3.2.1 Hard Constraints

- **Vehicle Dynamics**: The vehicle's future states must precisely follow the dynamics model. For each time step k:

$$x_{k+1} = f(x_k, u_k) \cdot \Delta t + x_k \tag{3.15}$$

  where $f(\cdot)$ is the nonlinear dynamics function.

- **Initial Condition**: The initial state of the prediction horizon is constrained to coincide with the currently observed vehicle state:

$$x_0 = x_{\text{current}} \tag{3.16}$$

- **State Bounds**: All state variables must remain within predefined physical and operational limits for all $k \in [0, N]$:

$$x_{\min} \leq x_k \leq x_{\max} \tag{3.17}$$

- **Control Input Bounds**: Control inputs must be within the actuator's physical capabilities for all $k \in [0, N-1]$:

$$u_{\min} \leq u_k \leq u_{\max} \tag{3.18}$$

- **Lane Boundaries**: The ego vehicle, approximated as a rectangle, must remain entirely within the lane boundaries. This constraint is typically expressed

in the Cartesian coordinate system by converting the Frenet states to Cartesian coordinates for the vehicle's corner points (V_e). For the lateral position Y_{ego}, the condition is

$$Y_{\text{lane,min}} + d_{\text{safety}} \leq Y_{\text{ego}} \leq Y_{\text{lane,max}} - d_{\text{safety}} \tag{3.19}$$

where a small safety margin d_{safety} is introduced to ensure adequate clearance from the lane markings.

3.3.2.2 Soft Constraints

Collision avoidance is formulated through soft constraints employing separating hyperplanes and slack variables, thereby ensuring optimization feasibility while imposing substantial penalties for constraint violations.

- **Separating Hyperplane Formulation**: For each agents' FRS and ego vehicle pair, possible collision avoidance is enforced through dynamically computed separating hyperplanes. Specifically, for agents' FRS i and prediction step k, a separating hyperplane is characterized by its normal vector $(-\lambda_{i,k}, 1)$ and intercept parameter $\mu_{i,k}$. The constraints ensure spatial separation between the ego vehicle and agents' FRS while permitting controlled violations through non-negative slack variables $S_{i,k}$:

$$\begin{bmatrix} -\lambda_{i,k} & 1 \end{bmatrix} \cdot V_{e,j,k} \geq \mu_{i,k} - S_{i,k} \qquad \forall j \in \mathcal{V}_{\text{ego}} \tag{3.20}$$

$$\begin{bmatrix} -\lambda_{i,k} & 1 \end{bmatrix} \cdot O_{i,m,k} \leq \mu_{i,k} \qquad \forall m \in \mathcal{V}_i \tag{3.21}$$

  where $V_{e,j,k} \in \mathbb{R}^2$ denotes the Cartesian coordinates of ego vehicle corner j at time step k, $O_{i,m,k} \in \mathbb{R}^2$ represents the coordinates of agents' FRS i's corner m at time step k, $\mathcal{V}_{\text{ego}}$ is the set of ego vehicle corner indices, and $\mathcal{V}_i$ is the set of corner indices for obstacle i. The slack variables $S_{i,k} \geq 0$ are subject to substantial quadratic penalties in the objective function to discourage constraint violations while maintaining optimization feasibility (detailed in Section 3.3.1).

- **Hyperplane Parameter Regularization**: To ensure numerical stability and prevent abrupt changes in the separating hyperplane orientation, the slope parameter $\lambda_{i,k}$ is constrained within a prescribed interval:

$$\lambda_{\min} \leq \lambda_{i,k} \leq \lambda_{\max} \tag{3.22}$$

  This regularization prevents ill-conditioned optimization problems and ensures smooth evolution of the hyperplane parameters across prediction horizons.

3.3.3 Simulation Parameters

The key parameters employed in the RMPC formulation are presented Table.3.1. These values were determined through theoretical analysis, established practices in autonomous vehicle control, and systematic empirical validation to ensure robust and effective system performance.

Table 3.1: Key Simulation Parameters

Parameter Category	**Parameter**	**Value**
MPC Configuration	Prediction Horizon (N)	30 steps
	Prediction Time (t_f)	3.0 s
	Discretization Interval (Δt)	0.1 s
Vehicle Specifications	Vehicle Length	4.6 m
	Vehicle Width	1.86 m
	Wheelbase	2.8 m
	Lane Width (W_{lane})	3.5 m
State Constraints	Arc Length (s)	$[0, 200]$ m
	Lateral Deviation (d)	$[-0.5W_{\text{lane}}, 0.5W_{\text{lane}}]$ m
	Heading Error (θ_e)	$[\frac{-\pi}{8}, \frac{\pi}{8}]$ rad
	Longitudinal Velocity (v)	$[0, 80]$ km/h
	Steering Angle (δ)	$[\frac{-40\pi}{180}, \frac{40\pi}{180}]$ rad
	Longitudinal Acceleration (a)	$[-6, 3]$ m/s^2
Control Constraints	Jerk (j)	$[-5, 5]$ m/s^3
	Steering Rate (ω)	$[\frac{-5\pi}{180}, \frac{5\pi}{180}]$ rad/s
State Penalty Weights	Lateral Deviation (w_d)	10.0
	Heading Error (w_{θ_e})	5.0
	Velocity Tracking (w_v)	1.0
	Steering Angle (w_δ)	0.2
	Acceleration (w_a)	0.5
Control Penalty Weights	Jerk Regularization (w_j)	100.0
	Steering Rate (w_ω)	200.0
Hyperplane Parameters	Minimum Slope ($\lambda_{\min}$)	-0.8
	Maximum Slope ($\lambda_{\max}$)	0.8
Slack Variable Penalties	Vehicle FRS	1×10^4
	Pedestrian FRS	1×10^3
	Terminal Progress (w_{term})	1×10^3
Safety Parameters	d_{safe}	0.2 m
	Minimum Vehicle Length	3 m

3.4 Evaluation Metrics

To provide a comprehensive and quantitative assessment of the proposed trajectory planning framework, we evaluate its performance against baseline methodologies using three fundamental criteria that collectively address the core requirements of autonomous driving systems: safety, efficiency, and comfort. These evaluation dimensions are specifically designed to capture the critical trade-offs inherent in navigating uncertain and occluded urban environments, where robust decision-making under partial observability is paramount.

3.4.1 Safety: Collision Rate

Safety represents the most critical performance criterion for any autonomous driving system, particularly in scenarios involving occlusions and uncertainty. We quantify safety performance through the **collision rate** metric, which serves as an absolute measure of system reliability.

Definition: The collision rate is defined as the percentage of experimental trials in which the ego vehicle's geometric footprint intersects at any point in time with that of any other road user (including vehicles or pedestrians emerging from occluded regions) or static environmental obstacles [12].

Mathematical Formulation:

$$\text{Collision Rate} = \frac{\text{Number of Trials with Collision}}{\text{Total Number of Trials}} \times 100\% \tag{3.23}$$

where a collision is detected when the geometric overlap between the ego vehicle and any other entity is non-zero at any time instant during the trial.

Significance: The fundamental objective of any motion planning algorithm, particularly robust controllers designed for occluded environments, is to achieve a collision rate approaching zero. This metric serves as a definitive, non-negotiable indicator of the planner's capability to guarantee safety-critical operation, even under challenging worst-case scenarios involving unpredictable road users emerging from blind spots. Any non-zero collision rate indicates a critical failure in the safety assurance mechanism.

3.4.2 Efficiency: Traversal Time

While maintaining absolute safety, an effective autonomous driving system must also demonstrate reasonable progress toward navigation objectives. We assess this performance dimension through the **traversal time** metric, which quantifies the temporal efficiency of the planning algorithm.

Definition: Traversal time, also referred to as task completion time or mission duration, represents the total temporal duration required for the ego vehicle to navigate from its designated initial position to the specified goal destination while adhering to traffic regulations and safety constraints.

Mathematical Formulation:

$$T_{traversal} = t_{goal} - t_{start} \tag{3.24}$$

where t_{start} and t_{goal} represent the timestamps of mission initiation and successful completion, respectively.

Significance: This metric is crucial for quantifying the "cost of robustness" inherent in conservative planning approaches [13]. A trajectory planner that exhibits excessive conservatism in the presence of occlusions may demonstrate significantly prolonged traversal times due to unnecessary deceleration maneuvers or extended waiting periods at potentially hazardous locations. Therefore, traversal time enables evaluation of the planner's capability to balance proactive safety measures with the maintenance of efficient traffic flow and reasonable progress rates. Lower traversal times indicate superior efficiency, provided that safety guarantees remain uncompromised.

3.4.3 Comfort: Root Mean Square Jerk

Passenger comfort constitutes a critical factor for the widespread acceptance and adoption of autonomous vehicle technology in practice. We quantify ride comfort through the analysis of motion **jerk**, which characterizes the smoothness and predictability of vehicle trajectories.

Definition: Jerk is defined as the time derivative of acceleration, mathematically expressed as $j(t) = \frac{da(t)}{dt}$, where $a(t)$ represents the vehicle's acceleration profile. Lower jerk magnitudes correspond to smoother motion characteristics with reduced abrupt changes in vehicle dynamics, thereby enhancing passenger comfort.

Mathematical Formulation: To evaluate the overall motion smoothness across the complete trajectory, we compute the Root Mean Square (RMS) of both longitudinal jerk (j_x) and lateral jerk (j_y) components over the entire traversal duration T:

$$J_{RMS,x} = \sqrt{\frac{1}{T}\int_0^T j_x(t)^2\,dt} \tag{3.25}$$

$$J_{RMS,y} = \sqrt{\frac{1}{T}\int_0^T j_y(t)^2\,dt} \tag{3.26}$$

The overall comfort metric is then computed as:

$$J_{RMS} = \sqrt{J_{RMS,x}^2 + J_{RMS,y}^2} \tag{3.27}$$

Significance: Aggressive maneuvering behaviors, such as emergency braking or sharp steering inputs when reacting to potential threats emerging from occluded

regions, typically result in elevated jerk magnitudes that degrade passenger comfort. This metric effectively captures the smoothness and predictability characteristics of the generated trajectories. Lower RMS jerk values are strongly preferred as they indicate superior ride quality and enhanced passenger acceptance, while also reducing mechanical stress on vehicle components and minimizing motion-induced discomfort [14].

The integration of these three complementary evaluation metrics provides a holistic assessment framework that enables comprehensive comparison of different trajectory planning approaches while highlighting the fundamental trade-offs between safety, efficiency, and comfort in autonomous driving systems operating under uncertainty.

4

Results

We conducted comprehensive experimental evaluations to assess the performance of our proposed occlusion-aware RMPC framework against a baseline method across two critical safety scenarios: hidden pedestrians and hidden vehicles. Each scenario was subjected to 200 independent simulation trials to ensure statistical significance. The baseline method employed a path-velocity decomposition approach for trajectory generation without occlusion-aware risk assessment. The reference path followed the designated route of the ego vehicle, with a predefined velocity profile that maintained normal driving speeds, applied moderate braking when approaching occluded regions, and resumed normal speeds once the occlusions were cleared.

4.1 Scenario 1: Hidden Pedestrian

4.1.1 Experimental Setup

In this scenario, we simulated the presence of a pedestrian randomly positioned within an occluded area, initially unobserved by the ego vehicle's sensor suite. The pedestrian was programmed to emerge onto the roadway at a random time during the ego vehicle's approach, with crossing directions uniformly distributed within $[-\frac{\pi}{8}, \frac{\pi}{8}]$ radians and velocities randomly sampled from $[1.0, 1.5]$ m/s, representing typical human walking speeds.

4.1.2 Qualitative Analysis

Figure 4.1 presents three representative snapshots at time instances $t = 3.5$s, 4.5s, and 5.3s, comparing the baseline method with our proposed approach. In this particular scenario, the ego vehicle traveled on the upper lane at 50 km/h while three red vehicles were parked on the adjacent lower lane, creating significant visual occlusions. A pedestrian positioned in the blind spot adjacent to the second parked vehicle, barely observable by the ego vehicle's sensors, began crossing the road perpendicularly at $t = 3.5$s with a velocity of 1.2 m/s.

The baseline method maintained a reference path along the road centerline with a velocity profile following a hyperbolic tangent curve, transitioning smoothly from 50 km/h to 30 km/h approximately 10 meters before the first parked vehicle, and recovering to 50 km/h when driving 5 meters beyond the third vehicle. As illustrated

in Figure 4.2 (row 1, column 1), this velocity profile resulted in the ego vehicle consistently following the centerline with modest speed reduction in occluded areas, ultimately leading to a collision with the pedestrian at $t = 5.3$s (Figures 4.1a, 4.1c, and 4.1e).

In contrast, our proposed method demonstrated proactive safety measures from the early stages of the scenario. As shown in Figure 4.1b, the ego vehicle initiated a slight lateral displacement away from the road centerline well in advance in anticipation of the occluded region. This strategic maneuver served dual purposes: first, it actively expanded the sensor field-of-view, enabling earlier detection of potentially hidden pedestrians; second, it increased the spatial separation from high-risk occluded areas. Upon entering the region where hidden pedestrians might emerge, the vehicle executed additional deceleration while maintaining the lateral offset. Consequently, as demonstrated in Figure 4.1f, when the pedestrian became visible, the ego vehicle maintained sufficient separation distance to avoid collision without requiring emergency maneuvers. As the occlusion cleared, the vehicle resumed acceleration and gradually returned toward the road centerline.

A critical aspect of our approach is the dynamic evolution of the reachable sets (depicted as pink areas in the figures). At $t = 3.5$s, as the ego vehicle approached the occluded region, reachable sets appeared adjacent to each parked vehicle, representing potential pedestrian emergence zones. The reachable set near the first vehicle was smallest due to the ego vehicle's proximity and reduced time-to-collision, while those near the second and third vehicles progressively increased in size. At $t = 4.5$s, after passing the first parked vehicle, its associated reachable set disappeared as the occlusion was resolved. The reachable set near the second vehicle became minimal due to the ego vehicle's proximity. By $t = 5.3$s, only the third vehicle's reachable set remained, which was significantly reduced due to the ego vehicle's acceleration and the shortened time until occlusion clearance. This dynamic reachable set evolution provides formal safety guarantees for the trajectory planning process.

Figure 4.2 provides a comprehensive comparison of system signals throughout the simulation, including velocity, acceleration, yaw angle, steering angle, and distance to the pedestrian. Our method demonstrated more pronounced deceleration and active steering inputs to minimize occluded areas. In the distance-to-pedestrian plot, the pedestrian's position before $t = 3.5$s is indicated by dashed lines (representing the occluded state). Our method maintained consistent separation from the pedestrian throughout the scenario, while the baseline method resulted in collision at $t = 5.3$s.

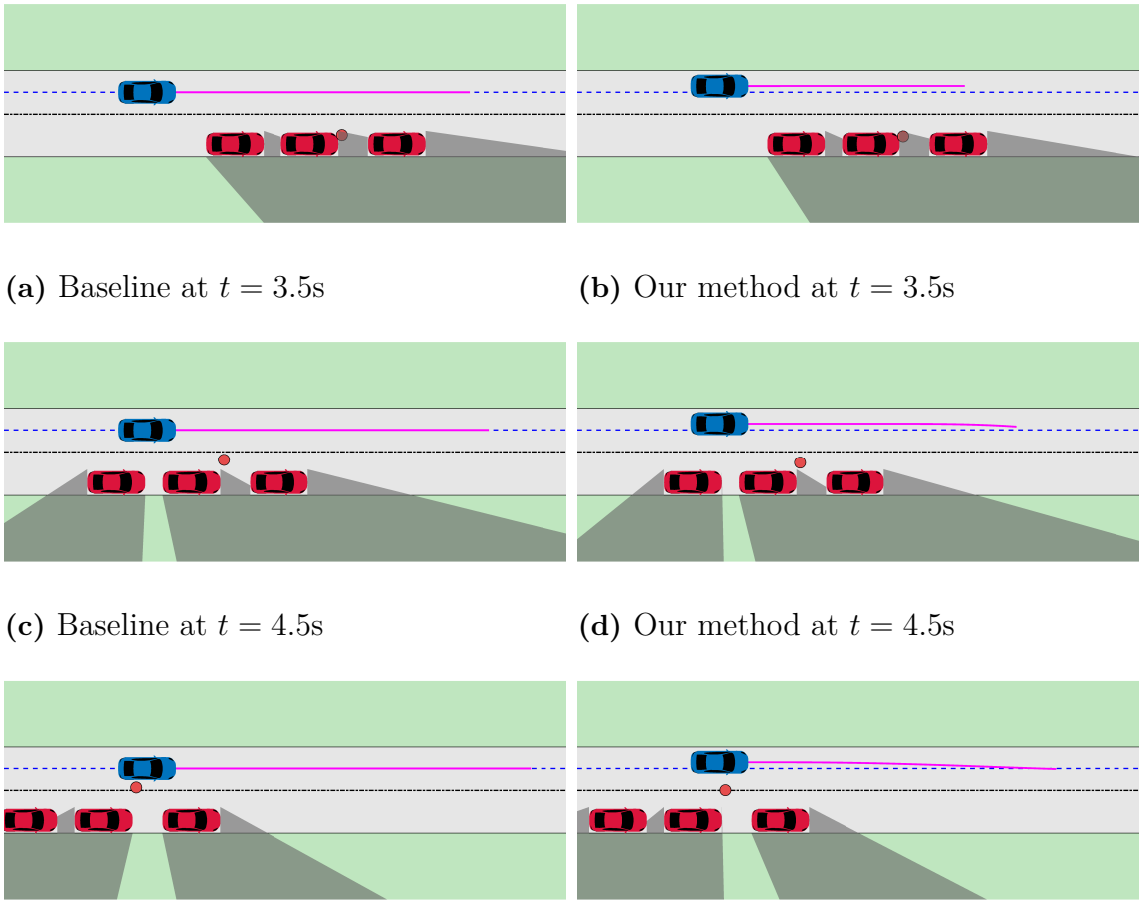

(a) Baseline at $t = 3.5$s

(b) Our method at $t = 3.5$s

(c) Baseline at $t = 4.5$s

(d) Our method at $t = 4.5$s

(e) Baseline at $t = 5.3$s

(f) Our method at $t = 5.3$s

Figure 4.1: Temporal evolution comparison between baseline method (left column) and our proposed approach (right column) at critical time instances: $t = 3.5$s, $t = 4.5$s, and $t = 5.3$s in the hidden pedestrian scenario.

4.1.3 Quantitative Results

The statistical results from 200 independent trials are presented in Table 4.1. Our proposed method achieved a perfect safety record with 0% collision rate, while the baseline method resulted in 17% collision rate. This safety improvement came at the cost of increased traversal time (7.2s vs. 6.5s, representing a 0.7s increase) and higher RMS jerk (2.2 m/s^3 vs. 1.3 m/s^3). The trade-off demonstrates that our method prioritizes safety while maintaining acceptable efficiency and comfort levels.

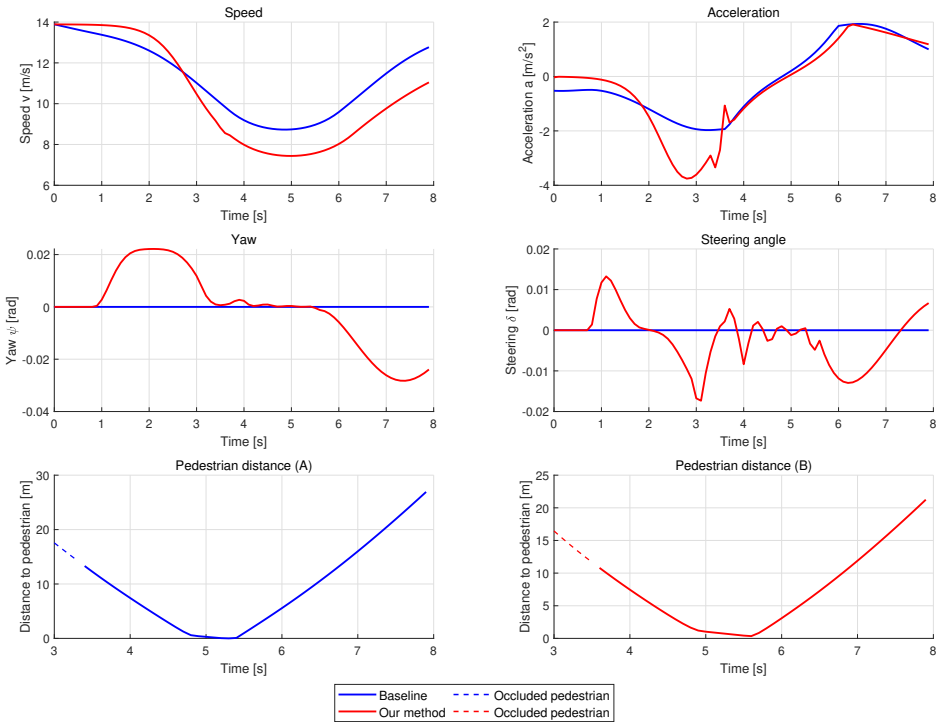

Figure 4.2: Comprehensive signal comparison for the hidden pedestrian scenario, showing velocity, acceleration, yaw angle, steering angle, and distance to pedestrian over time.

Table 4.1: Quantitative performance comparison between the proposed occlusion-aware RMPC and baseline method in Scenario 1 (Hidden Pedestrian).

Evaluation Metric	**Baseline**	**Our Method**
Collision Rate (%)	17.0	**0.0**
Traversal Time (s)	**6.5**	7.2
RMS Jerk (m/s^3)	**1.3**	2.2

4.2 Scenario 2: Hidden Vehicle

4.2.1 Experimental Setup

This scenario was designed around a two-way, single-lane intersection where the ego vehicle approaches with the intention of executing a left turn maneuver. Buildings positioned on both sides of the intersection create significant line-of-sight occlusions toward oncoming cross-traffic. The simulation incorporated at least one vehicle in the cross-traffic lane, with its velocity randomly sampled from 80% to 120% of the speed limit ($[0.8v_{\text{limit}}, 1.2v_{\text{limit}}]$) to represent realistic traffic variations.

4.2.2 Qualitative Analysis

Figure 4.3 presents a sequential analysis of the critical differences between our proposed method and the baseline approach when encountering vehicular occlusions. At $t = 4$s, a red cross-traffic vehicle was completely obscured by a large truck, as shown in Figure 4.3a (baseline) and Figure 4.3b (our method). Both control strategies initiated braking responses; however, the ego vehicle under our method's control applied significantly more pronounced deceleration.

By $t = 4.5$s, as illustrated in Figure 4.3d, the red vehicle became partially visible through the occlusion. Our method responded by further intensifying deceleration, demonstrating adaptive risk management based on evolving situational awareness. This resulted in the ego vehicle executing a smooth, controlled deceleration to a near-complete stop at the intersection by $t = 5$s (Figure 4.3f), thereby safely yielding right-of-way to the cross-traffic vehicle.

In stark contrast, the baseline method exhibited insufficient deceleration throughout the scenario, as demonstrated in Figures 4.3c and 4.3e. This inadequate response made collision with the red vehicle highly probable without emergency intervention maneuvers.

Figure 4.4 provides a comprehensive signal analysis comparing the proposed method with the baseline approach, illustrating the temporal evolution of key control and state variables throughout the intersection navigation scenario.

4.2.3 Quantitative Results

Figure 4.4 provides a comprehensive signal analysis comparing the proposed method with the baseline approach, illustrating the temporal evolution of key control and state variables throughout the intersection navigation scenario. Several critical behavioral differences are evident from this analysis:

When the ego vehicle approached the intersection with visual occlusions caused by roadside buildings, our method initiated deceleration significantly earlier than the baseline approach. A particularly challenging situation arose when a large truck appeared on the left lane, executing a right turn into the adjacent lane. This truck created substantial dynamic occlusion, obscuring potential cross-traffic vehicles. In

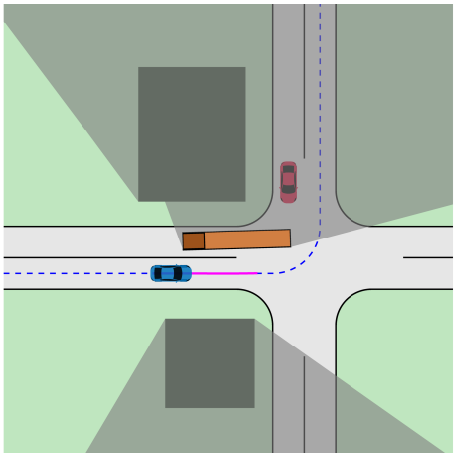

(a) Baseline at $t = 4$s

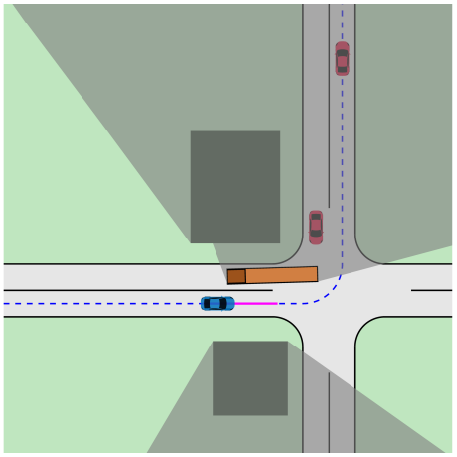

(b) Our method at $t = 4$s

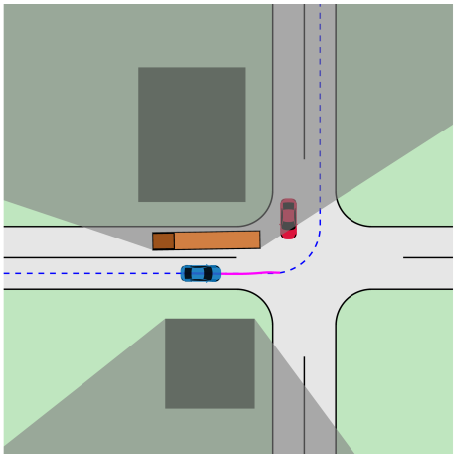

(c) Baseline at $t = 4.5$s

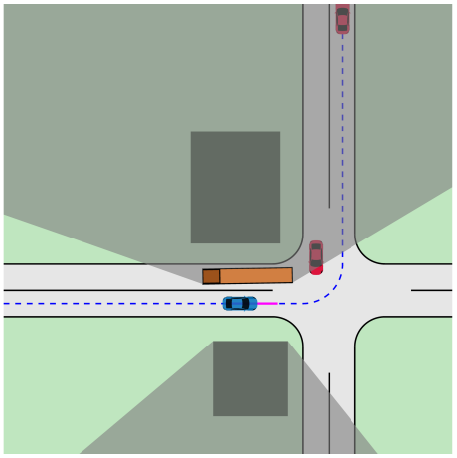

(d) Our method at $t = 4.5$s

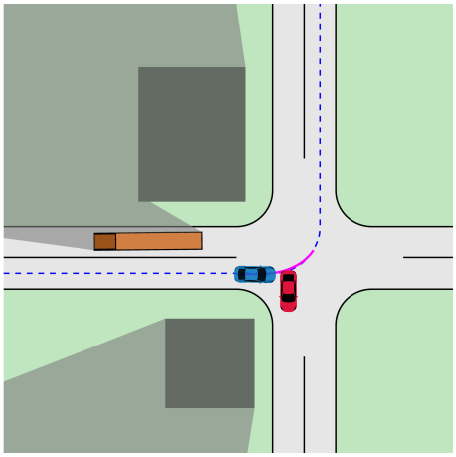

(e) Baseline at $t = 5$s

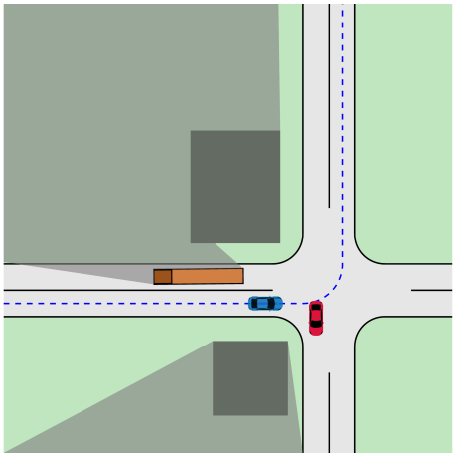

(f) Our method at $t = 5$s

Figure 4.3: Sequential comparison between baseline method (left column) and our proposed approach (right column) at critical time instances: $t = 4$s, $t = 4.5$s, and $t = 5$s in the hidden vehicle scenario.

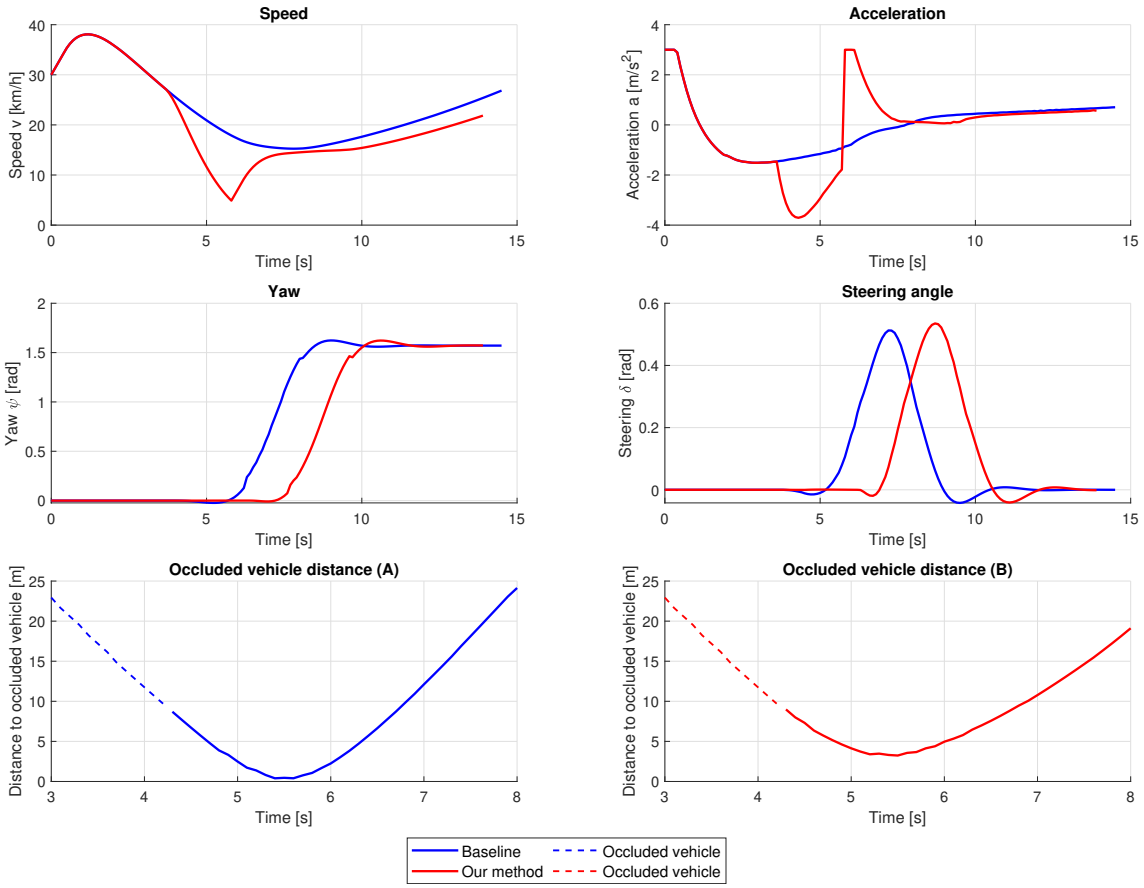

Figure 4.4: Comprehensive signal comparison for the hidden vehicle scenario, demonstrating the control strategy differences between baseline and proposed methods.

response to this compound occlusion scenario, our method decelerated the ego vehicle to a minimum speed of 5 km/h, maintaining this conservative velocity until the occluded regions were nearly completely resolved before initiating the left turn maneuver.

In contrast, the baseline method proceeded with the left turn prematurely, failing to adequately account for the dynamic occlusion created by the turning truck. This premature maneuver made collision avoidance impossible when the red cross-traffic vehicle emerged from behind the truck's occlusion zone. As demonstrated in the distance-to-vehicle comparison plot, the red vehicle became observable at approximately $t = 4.3$s. Throughout the scenario, our method maintained a safe separation distance from the emerging red vehicle, with the minimum separation distance never falling below 2.6 meters, thereby ensuring collision-free navigation.

The statistical analysis from 200 independent simulation trials is presented in Table 4.2. Our proposed method achieved perfect safety performance with 0% collision rate, while the baseline method resulted in 22.5% collision rate. Due to the more conservative deceleration strategy employed by our method, the traversal time increased from 10.7s to 12.5s, and the RMS jerk increased marginally from 2.1 m/s^3 to 2.8 m/s^3. These trade-offs demonstrate that our approach prioritizes collision avoidance while maintaining reasonable efficiency and passenger comfort.

Table 4.2: Quantitative performance comparison between the proposed occlusion-aware RMPC and baseline method in Scenario 2 (Hidden Vehicle).

Evaluation Metric	**Baseline**	**Our Method**
Collision Rate (%)	22.5	**0.0**
Traversal Time (s)	**10.7**	12.5
RMS Jerk (m/s^3)	**2.1**	2.8

4.3 Discussion

The experimental results demonstrate the effectiveness of our proposed occlusion-aware RMPC framework in achieving perfect safety performance across both challenging scenarios. The method successfully eliminated all collisions while maintaining reasonable efficiency and comfort trade-offs. The increased traversal times (10.8% in Scenario 1 and 16.8% in Scenario 2) and elevated jerk values represent the inherent cost of robust safety guarantees in uncertain environments. These results validate the fundamental design principle that proactive, conservative planning in occluded scenarios is essential for safe autonomous navigation in complex urban environments.

5

Conclusion

This chapter summarizes the key contributions of the thesis, discusses its limitations, and outlines potential directions for future research.

5.1 Summary

This thesis investigated the critical challenge of safe trajectory planning for autonomous vehicles operating in urban environments with significant sensor occlusions. A Robust Model Predictive Control framework was developed to explicitly account for potential hazards posed by unobserved road users. By incorporating formal reachability analysis, the framework estimates the spatio-temporal occupancy of hidden pedestrians and vehicles and integrates these estimates as soft constraints within the optimization problem via hyperplane separation. Comprehensive simulations of representative hidden pedestrian and hidden vehicle scenarios demonstrated the effectiveness of the proposed method. The RMPC planner achieved a zero-collision rate, representing a substantial improvement over a baseline MPC, which exhibited collision rates of 17.0% and 22.5%, respectively. This enhanced safety was attained at the cost of moderately increased traversal times and jerk, underscoring the inherent trade-off between robustness and performance. These results validate the RMPC approach as a viable strategy to ensure safety in uncertain, occluded driving conditions.

5.2 Limitations

Despite these promising results, several limitations should be acknowledged:

- **Simulation-only evaluation:** All experiments were performed in a simulated environment with idealized perception and simplified vehicle dynamics, introducing a potential sim-to-real gap [15].
- **Worst-case motion models:** The motion models for hidden agents rely on worst-case assumptions, which, while ensuring safety, can lead to overly conservative behavior in real-world interactive scenarios [16].
- **Fixed reference path:** The current framework assumes a fixed reference

path and reacts to occlusions rather than proactively maneuvering to improve visibility.

5.3 Future Work

Building on this work, several avenues for future investigation are recommended:

- **Real-world validation:** Progress from hardware-in-the-loop (HIL) simulations to full-scale on-road testing to bridge the sim-to-real gap and confirm the framework's practical applicability.
- **Hybrid prediction models:** Reduce conservatism by combining data-driven, probabilistic forecasts of agent behavior with the formal safety guarantees of reachability analysis.
- **Integrated path planning:** Couple the trajectory planner with a higher-level path planner to enable proactive information-gathering maneuvers, such as adjusting lateral position to gain improved visibility before blind corners.
- **Dynamic risk adaptation:** Enhance integration with upstream perception modules to dynamically adjust the controller's risk sensitivity based on real-time sensor confidence.

Bibliography

[1] Haoyang Fan et al. “Baidu apollo EM motion planner”. In: *arXiv preprint arXiv:1807.08048* (2018).

[2] Yihan Hu et al. “Planning-oriented autonomous driving”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2023, pp. 17853–17862.

[3] Jyh-Jing Hwang et al. “Emma: End-to-end multimodal model for autonomous driving”. In: *arXiv preprint arXiv:2410.23262* (2024).

[4] Chang Liu et al. “Path planning for autonomous vehicles using model predictive control”. In: *2017 IEEE Intelligent Vehicles Symposium (IV)* (2017), pp. 174–179. URL: https://api.semanticscholar.org/CorpusID:11835385.

[5] Ivo Batkovic et al. “Experimental Validation of Safe MPC for Autonomous Driving in Uncertain Environments”. In: *IEEE Transactions on Control Systems Technology* 31.5 (2023), pp. 2027–2042. DOI: 10.1109/TCST.2023.3291562.

[6] Ivo Batkovic et al. “Safe Trajectory Tracking in Uncertain Environments”. In: (2021). arXiv: 2001.11602 [eess.SY]. URL: https://arxiv.org/abs/2001.11602.

[7] Chris van der Ploeg et al. “Overcoming the Fear of the Dark: Occlusion-Aware Model-Predictive Planning for Automated Vehicles Using Risk Fields”. In: (2023). arXiv: 2309.15501 [cs.RO]. URL: https://arxiv.org/abs/2309.15501.

[8] Jiayu Fan, Nikolce Murgovski, and Jun Liang. “Efficient optimization-based trajectory planning”. In: (2024). arXiv: 2312.17440 [cs.RO]. URL: https://arxiv.org/abs/2312.17440.

[9] Hyunwoo Park et al. “Occlusion-aware risk assessment and driving strategy for autonomous vehicles using simplified reachability quantification”. In: *IEEE Robotics and Automation Letters* 8.12 (2023), pp. 8486–8493.

[10] Moritz Werling et al. “Optimal trajectory generation for dynamic street scenarios in a frenet frame”. In: *2010 IEEE international conference on robotics and automation*. IEEE. 2010, pp. 987–993.

[11] Markus Koschi and Matthias Althoff. “Set-based prediction of traffic participants considering occlusions and traffic rules”. In: *IEEE Transactions on Intelligent Vehicles* 6.2 (2020), pp. 249–265.

[12] Brian Paden et al. “A Survey of Motion Planning and Control Techniques for Self-Driving Urban Vehicles”. In: *IEEE Transactions on Intelligent Vehicles* 1.1 (2016), pp. 33–55. DOI: 10.1109/TIV.2016.2578706.

[13] Jaime F Fisac et al. “A general safety framework for learning-based control in uncertain robotic systems”. In: *IEEE Transactions on Automatic Control* 64.7 (2018), pp. 2737–2752.

[14] Rajesh Rajamani. *Vehicle dynamics and control*. Springer, 2006.

[15] Iván García Daza et al. “Sim-to-real transfer and reality gap modeling in model predictive control for autonomous driving”. In: *Applied Intelligence* 53.10 (2023), pp. 12719–12735.

[16] Xinwei Wang et al. “Reachability-based confidence-aware probabilistic collision detection in highway driving”. In: *Engineering* 33 (2024), pp. 90–107.

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY