Safe by Design Model Predictive Control for Autonomous Bus Driving

Public Report

Alejandro Blumentals
INRIA Grenoble

Pierre-Brice Wieber
INRIA Grenoble

December 13, 2017

This report describes the design of a safe trajectory generation module for an autonomous bus. In Section 1 we introduce the problem, its context and our goals. In Section 2 we describe the inputs and outputs of the trajectory generator. In Section 3 we assess the performance of our planner through simulations. We conclude the report with guidelines for future work in Section 4.

1 Our Goals

Context. Inria is now involved in the STAR project (Systeme de Transport Autonome Rapide), together with its industrial partners EasyMile, IVECO and Transdev. The industrial partners' ambitions are to develop an autonomous bus, which will be 12m long, will have a capacity of 110 people and will be able to drive at speeds higher than 40km/h while complying with passenger comfort and safety constraints. The buses will first operate in airports and in well controlled dedicated spaces. As the systems get better and the legislation evolves, the buses will drive in more complicated environments. Ultimately they will serve as public transportation systems in cities.

Main issue: Guaranteeing safety. An autonomous bus of 10 tons, driving at 40 km/h with 110 people standing in it, unattached, is a critical system. To commercialize an autonomous bus transportation system it is necessary to prove to the public and to the legislators that the system is safe. This requires having safe electronics, redundant sensors as well as a provably safe driving policy.

Our role: making driving policy safer. The driving policy of an autonomous vehicle is usually split into three parts: a behavioral layer, a trajectory generation layer and control layer [1]. Given our experience at Inria in the safe trajectory generation of robots, our role will be to contribute to the development of a safe trajectory generation module for the autonomous bus.

Current state of affairs and challenges ahead. The current autonomous shuttles developed by EasyMile and by other CityMobil2 project participants have been operating in their demonstrations at low speeds (10 - 20 km/h [2]). They have been successful in implementing a defensive driving policy that allows the vehicle to stop when pedestrians or cyclists cross its path. One of the challenges ahead is to design a driving policy which can scale up to cope with a larger vehicle driving at greater speeds and sharing the road with other fast-driving vehicles in a safe and comfortable way.

How to design a safe driving policy? The precise definition of what a safe driving policy should be, is still a matter of debate. In our opinion this debate is an important one and its resolution should come from a consensus between the users, the regulatory bodies and the AV technology providers. Lets now look at some examples of safety concepts and situations where they might succeed or fail.

Absolute safety. If the bus is parked and a cyclist runs into it then an accident cannot be avoided, *absolute safety* is not achievable, we cannot guarantee that accidents will never occur.

Passive safety. The next best thing which we can guarantee is that *if the bus has an accident then it will occur at zero speed*, this is known in robotics as *passive safety*. Our proposed trajectory generator implements this feature. Yet passive safety is not enough on its own, as it may lead to 0-speed accidents for which the bus may be held responsible. Passive safety means that the vehicle always has a maneuvre under the hood which will lead it to a full stop before a collision occurs. If the bus comes to a full stop on the wrong side of the road, leading to an accident, then the bus could considered responsible for the accident.

Responsibility Sensitive Safety (RSS). A common sense rule that the autonomous bus should enforce is to *never perform maneuvres which would lead to an accident of its own fault.* Clear rules assigning fault in an accident involving an autonomous bus should therefore be formulated and implemented by the behavioral layer of the bus. A good starting point for this would be to adapt the model proposed in [3], called *Responsibility Sensitive Safety (RSS)*, for the particular settings in which the autonomous bus will drive. Yet even RSS and passive safety may still be insufficient. If the bus is at a halt in its lane and another vehicle unexpectedly starts backing up towards the bus, then a human bus driver would probably honk and back up thus avoiding the collision, instead of colliding at zero speed and resting assured that it was the other vehicle's fault. In this case a suboptimal behavior of the autonomous bus may be partly blamed for the accident.

The conclusion is that design of a safe driving policy requires a safe behavioral module and a safe trajectory generation module. While the definition of safety itself is still evolving we propose that the trajectory generator should at least comply with passive safety.

In the next section we describe the design of our trajectory generation module.

2 Trajectory generation module

We wish to tackle the trajectory generation problem for an autonomous bus. The trajectory generation module is one of the main components required for the autonomous navigation of a vehicle [1]. Its responsibility is to take as input the desired driving behavior of the vehicle, information about the vehicle's state and of its environment and to output a trajectory. The output trajectory or motion plan is then to be tracked by a local feedback controller. In the following we describe more in detail the assumptions and constraints on the inputs and outputs of the module.

2.1 Inputs of the module

Desired driving behavior. The input desired driving behavior could be a lot of things, for example it could be 'stay on the current lane', 'change lanes', 'advance on at some desired speed', 'pass through some checkpoints' or 'stop at that point'. We assume that this information is provided to the motion planning module by a behavioral decision layer which is capable of making decisions such as negotiating an intersection or knowing that the vehicle has to stop at a red light. Essentially the role of the behavioral layer is to set up the correct tasks for the trajectory generator.

The exact form in which the desired driving behavior is communicated to our module is still an open question on our end. This should be further discussed with the party developing the behavioral layer for the bus.

For the moment our implementation assumes that a reference trajectory in position, velocity and heading is given. The reference trajectory could model the centerline of the virtual lane on which the bus drives. Implicit in the desired driving behavior is that the bus should avoid collisions with static and dynamic obstacles. Also a comfort value for the acceleration and jerk, as well as the minimum and maximum speed should be provided as part of the desired driving behavior.

State of the vehicle and of the environment. The input state of the vehicle is the result of some state estimation procedure which relies on noisy data and on an inexact knowledge of the environment. We must thus assume that the state of the vehicle and of its environment is not known exactly and the planned motion must then be robust with respect to this uncertainty.

The following table summarizes the inputs of the trajectory generation module.

trajectory generator Input	Symbol
Current vehicle state	$\mathbf{X_0}$
Current Plan	$\tilde{\mathbf{U}}$
Static Obstacles	$\mathcal{O}_s$
Dynamic Obstacles	$\mathcal{O}_d(t)$
Maximum Speed	v_{max}
Minimum Speed	v_{min}
Maximum Acceleration	a_{max}
Maximum Longitudinal jerk	$\dot{a}_{max}$
Maximum steering angle	φ_{max}
Maximum steering rate	$\dot{\varphi}_{max}$
Desired trajectory	$\mathbf{x^{ref}(t)}$
Desired speed	$v^{ref}(t)$
Desired heading	$\mathbf{d_1{}^{ref}(t)}$
Comfort acceleration	a_{comf}
Comfort jerk	$\dot{a}_{comf}$

2.2 Output Trajectory.

The planned motion must comply with several criteria.

1 **Kinematic feasibility.** The path must be kinematically feasible, this includes the non-holonomic constraints as well as the limitations of the steering angle and on the acceleration capacity of the bus.

2 **Safety of the passengers.**

- The bus must avoid collisions with static and dynamic obstacles.
- The bus must be able to execute at all times an emergency maneuvre which leads it to a full stop before any collision happens.
- The bus must obey traffic laws such as the maximum speed limit.

3 **Comfort of the passengers.** The driving should feel smooth and comfortable, so no sudden changes in velocity, acceleration and steering angle should arise. Standing passengers shouldn't feel that they're losing their stability. The motion of the bus must be such that the passengers don't fall.

4 **Objectives.** The bus must comply with the desired driving behavior.

2.3 Vehicle specifications

Further parameter values which characterize the autonomous bus are summarized in the following table.

Quantity	Symbol	Value
Mass	m	$20 \cdot 10^3$ kg
Length	L	12 m
Inter Axle Distance	l	10 m
Width	w	2.55 m
Maximum Operating Speed	V_{max}	40 km/h
Comfort acceleration	a_{comf}	$0.93ms^{-2}$
Comfort jerk	$\dot{a}_{comf}$	$0.60ms^{-3}$
Capacity of the bus	C	110 people

The comfort accelerations and jerks correspond to the maximum which a human can withstand without losing balance, according to [4]. It would be good to know the opinion of the bus constructors or operators on these values. Most of the studies we've come across concern only rail vehicles.

3 Current Results

3.1 Overtaking a small obstacle on a two lane road

Figure 1 depicts the results of a simulation where the vehicle must stay within a two lane road and track the centerline of the right lane at 40 km/h. A small static obstacle blocks the vehicle's path. The planned trajectory safely avoids the obstacle and stays within the lanes. After the obstacle is avoided, the vehicle merges back to right lane and keeps a 40 km/h velocity. We observe, however, that the merge takes about 200 metres to complete. The gains could be modified to enable a more accurate tracking of the trajectory and a smaller merge distance if needed. We can also note that the maximum acceleration and jerk on the bus exceed the comfort thresholds, suggesting that there is room for improvement in the handling of passenger comfort.

The computation time of solving each NLP was on average 17 ms on an Intel Core i7 2.6 GHz processor.

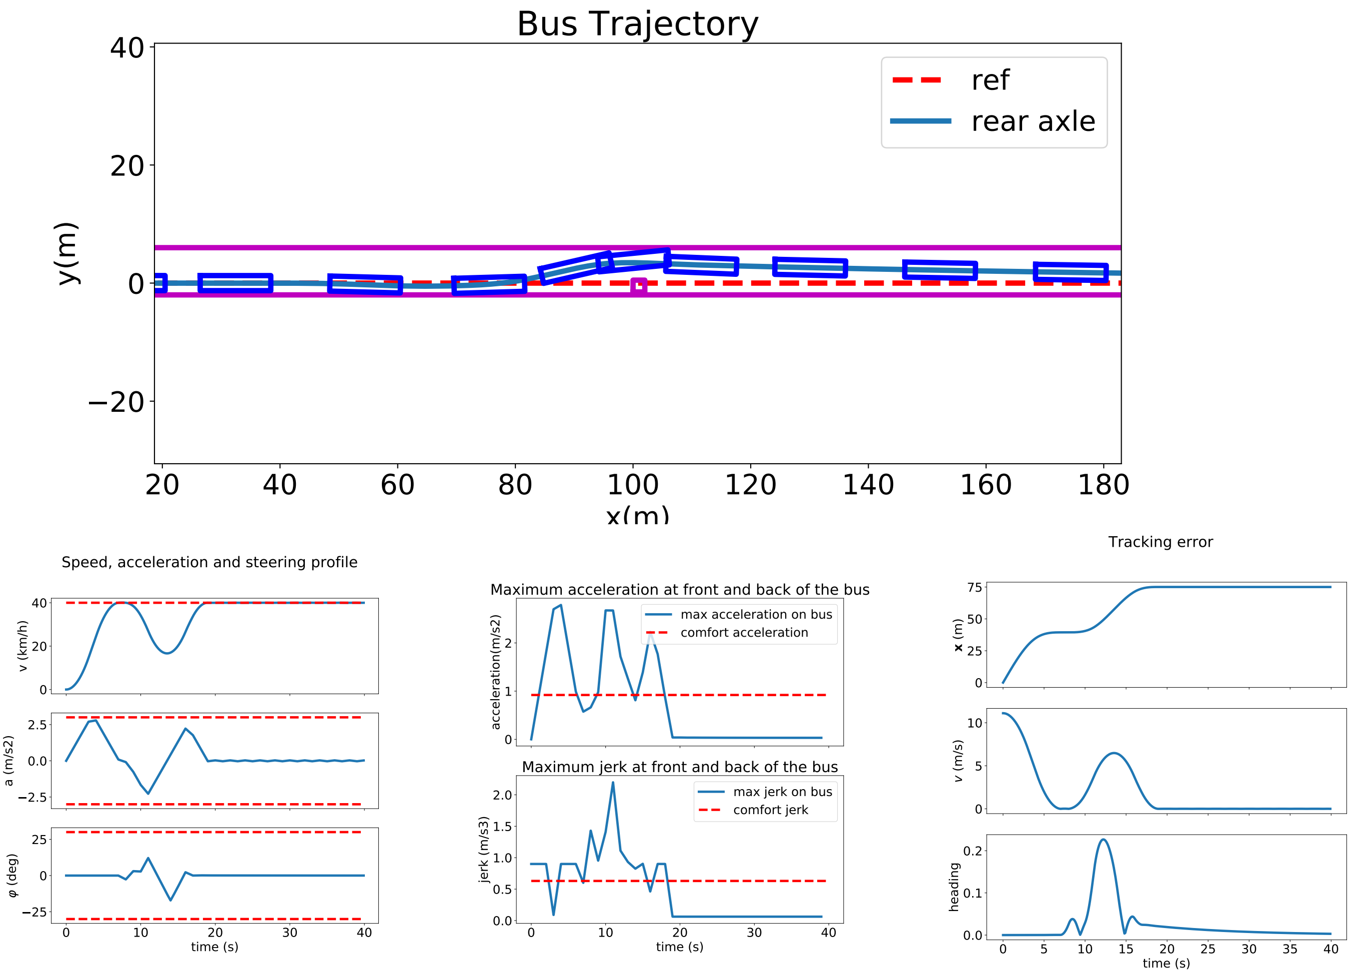

Figure 1: Overtaking a small obstacle

3.2 Trying to overtake a large obstacle on a two lane road

Figure 2 depicts the results of a simulation where the vehicle must stay within a two lane road and track the centerline of the right lane at 40 km/h. A large static obstacle blocks the vehicle's path such that overtaking is impossible. The planned trajectory safely stops before a collision with the obstacle or with the lane boundaries occurs. Although it is hard to see at the scale of the plots, an imposed safety distance of 10 cm is kept between the bus and the obstacles. This safety distance could be modified as needed. Again, the comfort of the maneuvre could still be improved.

The computation time of solving each NLP was on average 26 ms on an Intel Core i7 2.6 GHz processor.

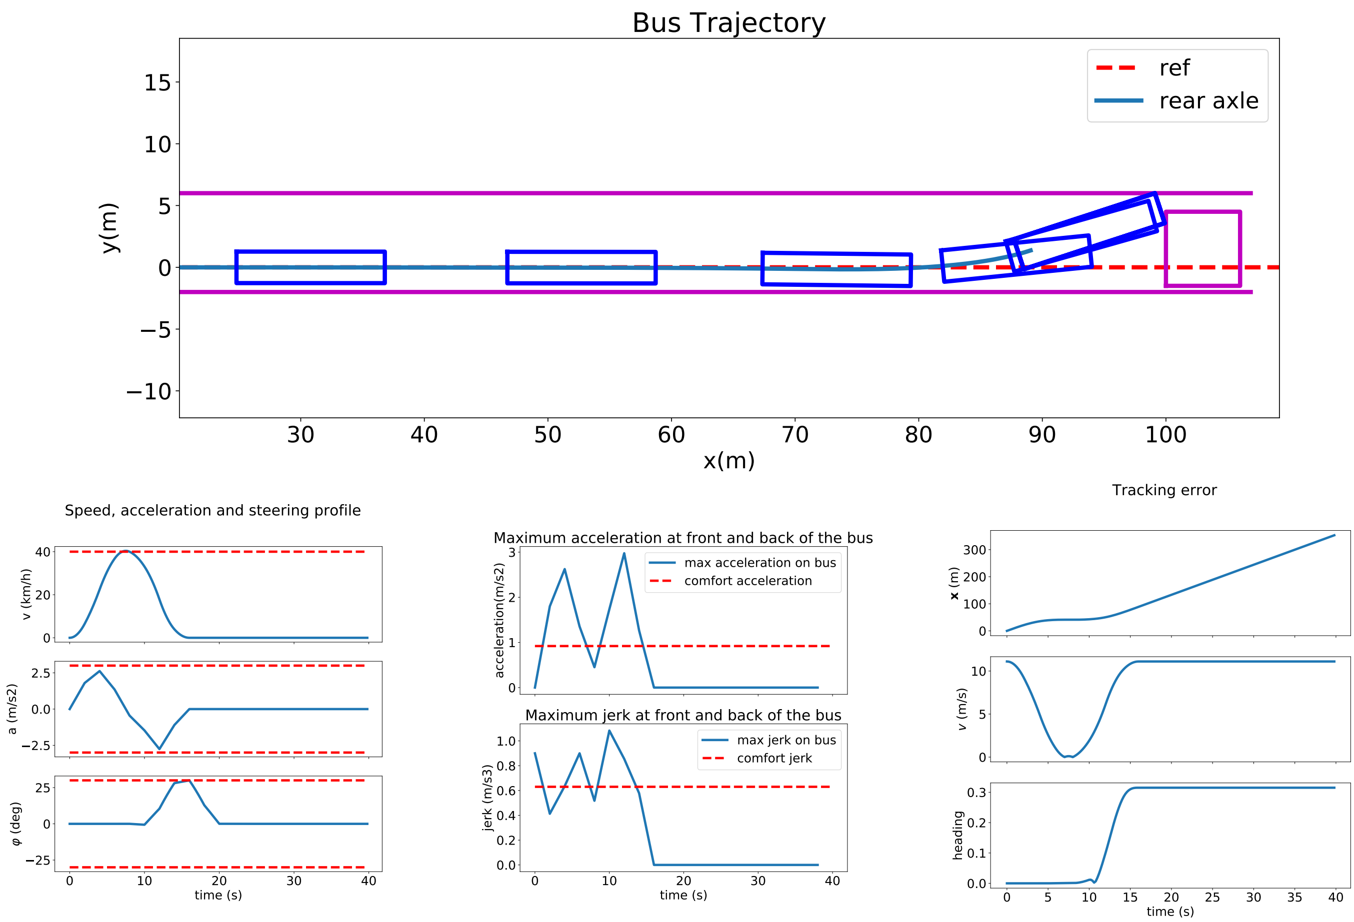

Figure 2: Trying to overtake a large obstacle

3.3 Double Lane change

Figure 3 depicts snapshots from a double lane change simulation as described in ISO 3888-2. There are two lanes, the vehicle starts on the right, it must change lanes to the left and then again to the right in a very short distance. In a real experiment cones are placed along the trajectory to define the drivable space and the actual physical vehicle has to avoid hitting the cones. In the simulation in Figure 3 the vehicle safely avoids the obstacles (depicted in magenta) while tracking the centerline of the drivable space at a reference speed of 40 km/h. The vehicle's planned trajectory at each of the instants is depicted in green.

Note however that the scenario in ISO 3888-2 was initially designed for cars, so the vehicle in this simulation is only 4 metres long. It would be a good idea to define such a standard maneuvre for a bus. We could then compare simulation results with human driving data and maybe even try to fit the parameters in our model so as to produce human-like driving maneuvres for the autonomous bus.

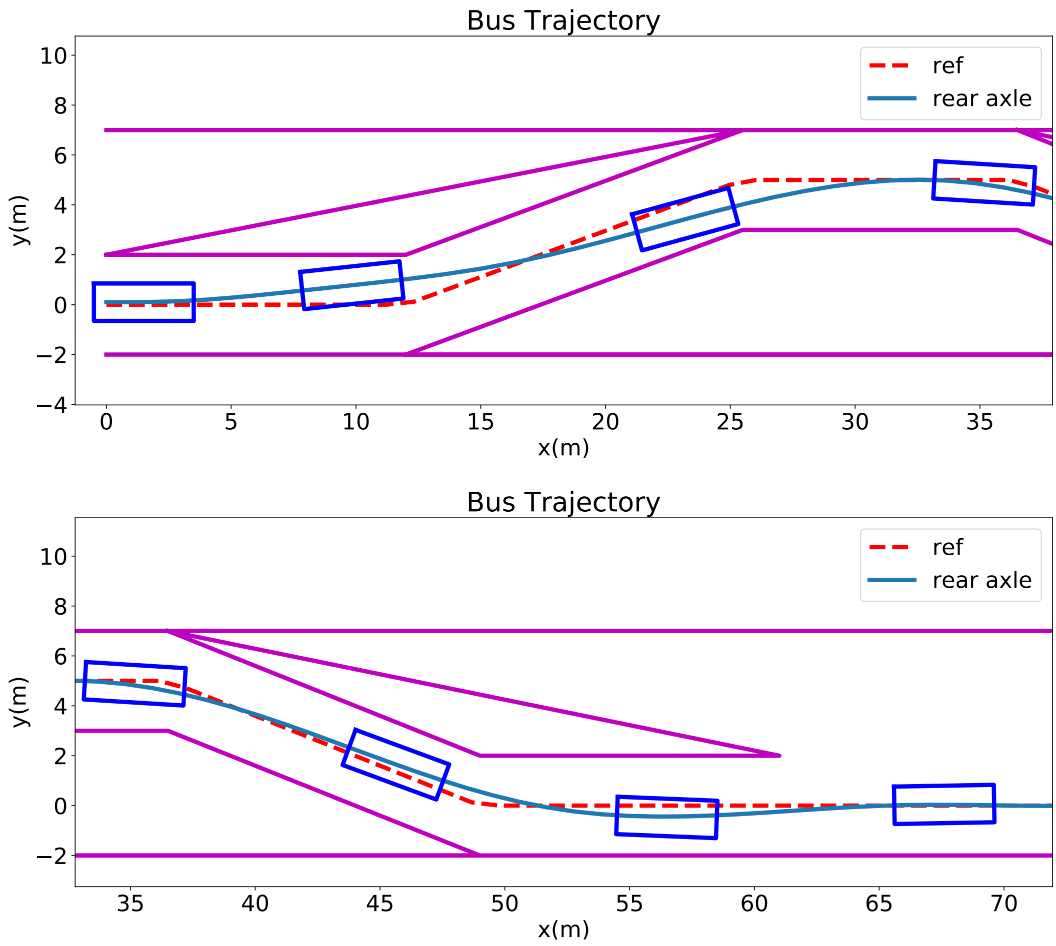

Figure 3: Double lane change

3.4 Double Lane change with a suboptimal behavior

Figure 4 depicts snapshots from the same simulation as in the previous section, but where we have replaced the obstacles by small triangles (modelling the cones which are placed in the physical experiment). This results in a trajectory where the vehicle stops before it hits a cone and then stays still. The trajectory is safe with respect to the given constraints yet it is suboptimal. Introducing 18 triangles as constraints has artificially created a local minimum where the trajectory generator gets stuck. The lesson to be learned is that the trajectory generator must be given input constraints which adequately represent the drivable space of the vehicle.

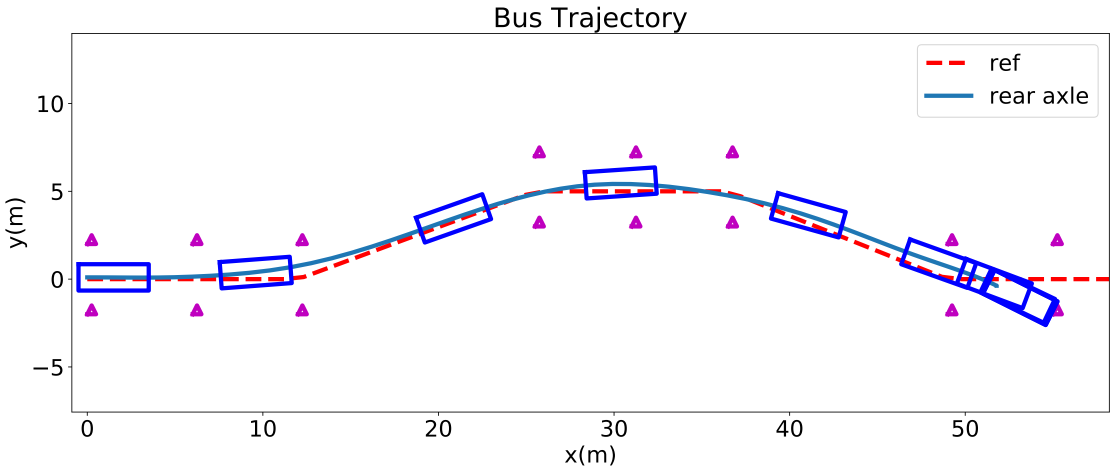

Figure 4: Double lane change with a suboptimal behavior

3.5 Avoiding obstacles in an unstructured environment

Figure 5 depicts the results of a simulation where the vehicle is asked to track a square trajectory at 40 km/h. One large static obstacle and two smaller ones block the vehicle's path. The planned trajectory safely avoids collision with the obstacles. Note that the reference path is kinematically infeasible for the vehicle because of the sharp corners, the generated trajectory smooths this out. The simulation also illustrates that there may exist many local minima to the problem, as obstacles may be avoided from the left or from the right. Our method finds local minima and hence the vehicle chooses to avoid obstacles to the left or to the right depending on which option changes its current plan the least.

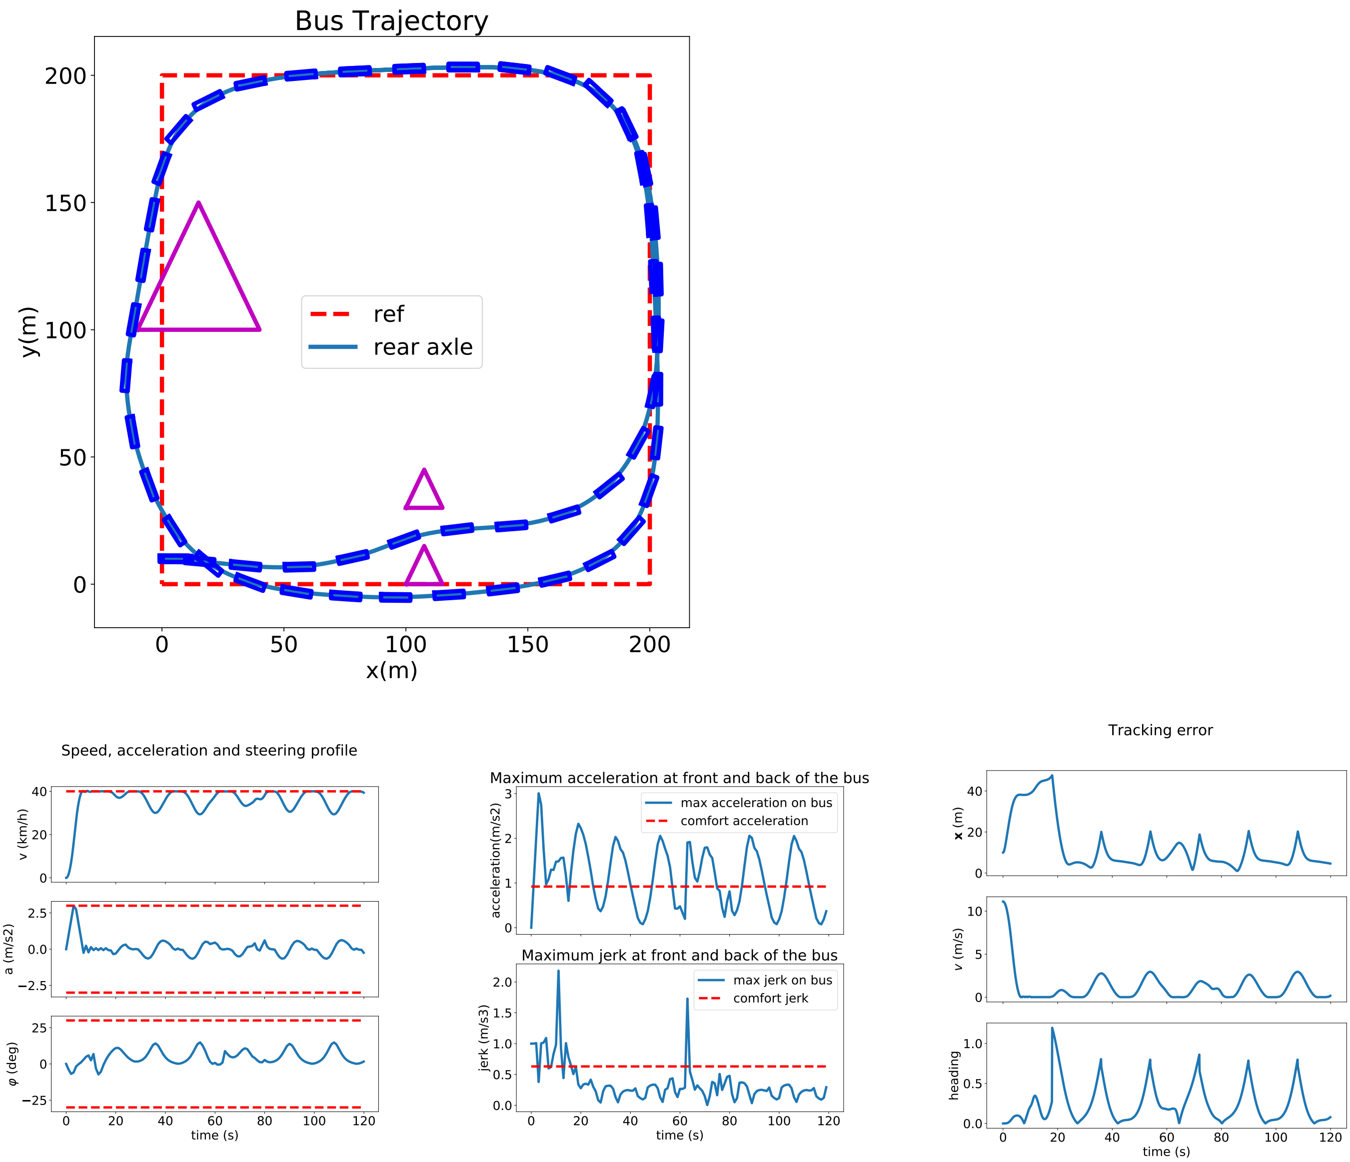

Figure 5: Driving in an unstructured environment

4 Current Conclusions

Recapitulation. We have presented the design of a trajectory generation module for an autonomous bus based on Model Predictive Control. The module implements passive safety, thus guaranteeing that the bus always comes to a full stop before any accident occurs. We demonstrated the capabilities of the trajectory generation module through simulations, where we showed that collision avoidance maneuvres could be safely planned and executed while driving at 40 km/h.

Limitations and future work.

- **Comfort.** Our simulations show that the planned maneuvres may still be too uncomfortable for the passengers of the bus. We believe that introducing virtual constraints to enforce using the comfort deceleration whenever possible would improve this issue, as suggested in Equation (**??**).
- **Dynamic obstacles.** Our prototype implementation only covers the case of static obstacles, dynamic obstacles must be included.
- **Safety of the computation.** For now we have focused on the safety of the result, that is if our NLP solver converges then it will yield a safe trajectory for the vehicle to follow. However if the solver fails to compute the solution at some point, then we must execute a safe maneuvre from the previous plan. It would be beneficial to solve the MPC problem in such a way that each iterate is guaranteed to respect the constraints, thus forcing the trajectory planner to *always* yield a safe maneuvre, even if it has not converged.
- **Interfacing with the behavioral layer.** To guide the future development of the trajectory generation module it will be important to take into account how it is expected to interact with the behavioral layer.

References

[1] B. Paden, M. Cap, S. Z. Yong, D. Yershov, and E. Frazzoli, "A Survey of Motion Planning and Control Techniques for Self-Driving Urban Vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, 2016.

[2] B. Pessario, "Evaluation of Automated Vehicle Technology for Transit," no. 2117, pp. 1–21, 2016.

[3] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "On a formal model of safe and scalable self-driving cars," *CoRR*, vol. abs/1708.06374, 2017.

[4] B. D. Graaf and W. V. Weperen, "The Retention of Blance: An Exploratory Study into the Limits of Acceleration the Human Body Can Withstand without Losing Equilibrium," *Human Factors*, vol. 39, no. 1, pp. 111–118, 1997.