Ensuring Motion Safety of Autonomous Vehicles through Online Fail-safe Verification

Christian Pek and Matthias Althoff
Cyber-Physical Systems Group, Technical University of Munich
Email: {christian.pek, althoff}@tum.de

I. Motivation

Safety is undoubtedly the most important factor for the success of autonomous vehicles [13]. However, ensuring their safety is a challenging task since they operate in highly uncertain environments with multiple dynamic obstacles whose future motions are unknown [25, 24]. Even when trying to accurately predict and consider the most likely trajectory of obstacles, planned motions might become unsafe when obstacles deviate from the prediction, which regularly happens in real traffic. As a result of the arising uncertainties, most motion planning methods cannot exclude the possibility that the autonomous vehicle causes a collision. For instance, a residual risk of 0.1% per journey can imply one collision in 1,000 journeys. In order to solve safe motion planning, novel approaches need to be developed which 1) provably guarantee safety even if obstacles suddenly deviate from the predicted behavior and 2) are able to cope with any new occurring traffic situation and measurement uncertainties on the fly.

II. Related Work

Many planning approaches compute trajectories which are collision-free against the predicted most likely trajectories of obstacles within the planning horizon of typically $5\,\mathrm{s}$ to $12\,\mathrm{s}$ [26, 27, 10]. In safety-critical situations, the approaches in [1, 6] plan evasive trajectories while accounting for likely collisions with obstacles. However, when only predicting single behaviors, the safety of planned trajectories only holds if all obstacles do not deviate from this behavior.

One way to guarantee the safety of the autonomous vehicle is logical reasoning. Here, planned motions are checked whether they comply with certain rules and axioms, e. g., formulated using higher-order logic [8]. For instance, the safety of planned lane-change maneuvers is assessed in [12] and the safety of vehicle following in [16]. Although logical reasoning guarantees safety, logical expressions used for the verification often become highly complex and must be adapted to new, previously unmodeled traffic scenarios.

Yet in other approaches [23, 17, 2], planned motions are only executed if they do not contain *inevitable collision states* (ICSs). ICSs are states from which the autonomous vehicle eventually collides, no matter what trajectory it executes. The less-restrictive concept of *passive safety* demands that the vehicle is at a standstill at the time of collision, ensured by evaluating pre-computed braking trajectories [7]. Both ICS and passive safety are computationally expensive, and most works are only online capable when considering a small number of predicted trajectories of obstacles.

Reachability analysis, in contrast, accounts for any feasible future motion of dynamic obstacles [18, 11, 3]. The reachable set of obstacles is the set of states reachable from an initial set of states considering all feasible trajectories. Future collisions are identified by checking the computed reachable sets for intersections with the occupancy of the autonomous vehicle in the position domain [9, 4]. However, set-based techniques have the disadvantage that unsafe regions may grow rapidly for long planning horizons, eventually blocking the whole drivable area of the autonomous vehicle. Thus, many long-term plans are rejected as being potentially unsafe.

III. Proposed Approach

While one cannot exclude that autonomous vehicles may be part of an accident, one can eliminate self-inflicted collisions [25]. Therefore, we propose to verify the safety of each planned motion of the autonomous vehicle on the fly before execution assuming that obstacles obey traffic rules with reasonable care. Thus, if a collision occurs, the autonomous vehicle is not responsible. In every planning step, we verify planned intended motions with the following two steps.

Using reachability analysis, we first compute possibly occupied regions in the environment by considering all feasible trajectories and measurement uncertainties of obstacles. Our obtained occupancy sets are over-approximative and thus capture all feasible legal behaviors of obstacles [14, 15]. In order to obtain tight over-approximations, we assume that obstacles adhere to traffic rules, e. g., respecting a certain speed limit. However, we can remove these constraints individually at any time, e. g., at latest when obstacles violate traffic rules.

In a second step, we compute fail-safe trajectories. These trajectories branch off at the intended motion of the autonomous vehicle, denoted as ego vehicle in the following, and are constrained to not intersect with any of the occupancy sets (cf. Fig. 1) and to end in a safe state. The combination of intended motion plans with fail-safe trajectories considers all possible behaviors of obstacles and thus ensures the safety of the ego vehicle. For instance, if obstacles deviate from their predicted most likely motion, the ego vehicle has two options to remain safe along its intended motion: (1) execute the previously computed fail-safe trajectory (cf. Fig. 1a), or (2) find a new combination of an intended motion and a fail-safe trajectory (cf. Fig. 1b). The previously computed fail-safe

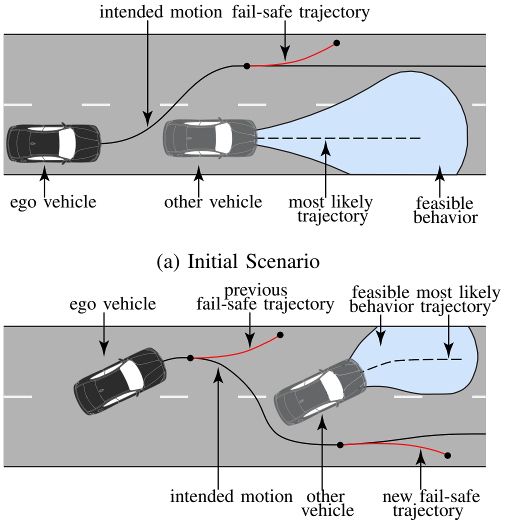

(a) Initial Scenario

(b) Future Scenario

Figure 1. Combining intended motion plans with fail-safe trajectories to guarantee safety. (a) Fail-safe trajectories are collision-free with respect to any feasible behavior of obstacles. (b) While the ego vehicle moves along its intended motion, new fail-safe trajectories are computed. If no valid new fail-safe trajectory is found, the ego vehicle must execute the previously computed fail-safe trajectory.

trajectory must be executed at the latest if the ego vehicle is located at the start of the fail-safe trajectory and no new fail-safe trajectory has been found after this point. Note that even though the ego vehicle had to start executing a fail-safe trajectory, it can recover and return to its comfort driving mode by verifying a new intended motion if the safety-critical situation resolves.

IV. Preliminary Results

We have already solved several challenges on the way to using our online verification approach in autonomous vehicles [22]. Our open-source tool SPOT predicts all feasible legal behaviors of obstacles in a rigorous and over-approximative way [14]. In order to obtain fail-safe trajectories, we first determine the *point of no return*, i. e., the last state along the intended motion for which a fail-safe trajectory still exists [21]. Secondly, we compute the fail-safe trajectory itself by solving convex trajectory optimization problems which consider the lateral and longitudinal dynamics of the vehicle separately [20]. Furthermore, we integrate safe states as desired goal states within the optimization to keep the ego vehicle safe for an infinite time horizon, i.e., never entering ICSs [19].

We have already validated our approach in simulation using a variety of recorded traffic scenarios [5]. Our approach verifies planned motions of the ego vehicle in less than $20\,\mathrm{ms}$ on average on a computer with an Intel i5 1.4GHz processor and 8 GB of DDR3 1600 MHz memory. Fig. 2 shows an urban T-junction scenario with three obstacles $b_i, i \in \{1, 2, 3\}$. Since the ego vehicle has the right of way, it plans a collision-free intended motion considering the most likely trajectories of all obstacles b_i (cf. Fig. 2a). However, if obstacle b_2 overlooks the ego vehicle and as a result disrespects its right of way, the intended motion ends in a collision. Based on the available free space, our approach ensures safety by computing a fail-safe trajectory (time horizon of $6\,\mathrm{s}$) which lets the ego vehicle turn right and come to a stop behind the occupancy set of b_3. This fail-safe trajectory starts at the last possible point in time along the intended motion (note that a braking maneuver without turning right needs to be executed earlier). Up until this point in time, the ego vehicle can already start to verify the newly planned intended motion while accounting for new sensor measurements.

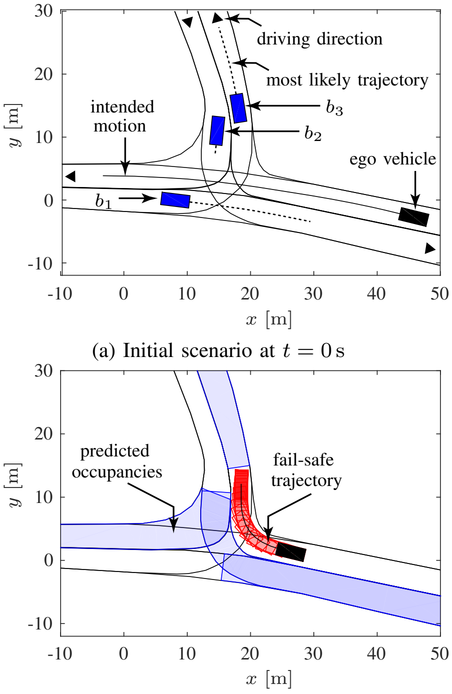

(a) Initial scenario at $t = 0\,\mathrm{s}$

(b) Predicted scenario at $t = 6\,\mathrm{s}$

Figure 2. Urban T-junction scenario (CommonRoad-ID: DEU_Ffb-2_2_S-1:2018b) in which vehicle b_2 violates the right of way rule of the ego vehicle. The predicted occupancies are only shown for $t = 6\,\mathrm{s}$ for clarity. The ego vehicle can avoid a collision by executing the computed fail-safe trajectory, which lets the ego vehicle turn right and come to a stop behind the occupancy set of obstacle b_3.

V. Conclusions and Future Work

Our proposed online verification approach guarantees the safety of autonomous vehicles on the fly. In every planning cycle, we verify intended motions by predicting all feasible legal behaviors of obstacles and computing fail-safe trajectories. In contrast to existing works, our approach is computationally efficient and enables fail-safe operation, since a safe plan exists at any given point in time—even if a newly planned intended motion of the vehicle is rejected as potentially unsafe.

Currently, we are testing our approach on real test vehicles to validate the safety benefits and to investigate the intervention rate of our safety layer in various traffic situations. Moreover, we are planning to further benchmark our approach in simulated safety-critical scenarios.

References

[1] C. Ackermann, J. Bechtloff, and R. Isermann. Collision avoidance with combined braking and steering. *6th Int. Munich Chassis Symposium*, pages 199–213, 2015.

[2] D. Althoff, M. Buss, A. Lawitzky, M. Werling, and D. Wollherr. On-line trajectory generation for safe and optimal vehicle motion planning. In *Autonomous Mobile Systems*, pages 99–107. 2012.

[3] M. Althoff. Reachability analysis of nonlinear systems using conservative polynomialization and non-convex sets. In *Proc. of Hybrid Systems: Computation and Control*, pages 173–182, 2013.

[4] M. Althoff and J. M. Dolan. Online verification of automated road vehicles using reachability analysis. *IEEE Transactions on Robotics*, 30(4):903–918, 2014.

[5] M. Althoff, M. Koschi, and S. Manzinger. CommonRoad: Composable benchmarks for motion planning on roads. In *Proc. of the IEEE Intelligent Vehicles Symposium*, pages 719–726, 2017.

[6] S. J. Anderson and S. C. Peters. An optimal-control-based framework for trajectory planning, threat assessment, and semi-autonomous control of passenger vehicles in hazard avoidance scenarios. *Int. Journal of Vehicle Autonomous Systems*, 8:190–216, 2010.

[7] S. Bouraine, T. Fraichard, and H. Salhi. Provably safe navigation for mobile robots with limited field-of-views in dynamic environments. *Autonomous Robots*, 32(3): 267–283, 2012.

[8] W. Damm, H.-J. Peter, J. Rakow, and B. Westphal. Can we build it: formal synthesis of control strategies for cooperative driver assistance systems. *Mathematical Structures in Computer Science*, 23(04):676–725, 2013.

[9] P. Falcone, M. Ali, and J. Sjöberg. Predictive threat assessment via reachability analysis and set invariance theory. *IEEE Transactions on Intelligent Transportation Systems*, 12(4):1352–1361, 2011.

[10] E. Frazzoli, M. A. Dahleh, and E. Feron. Real-time motion planning for agile autonomous vehicles. In *Proc. of the American Control Conference*, pages 43–49, 2001.

[11] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin. FaSTrack: A modular framework for fast and guaranteed safe motion planning. In *Proc. of the IEEE Conference on Decision and Control*, pages 1517–1522, 2017.

[12] M. Hilscher, S. Linker, and E.-R. Olderog. Proving safety of traffic manoeuvres on country roads. In *Theories of Programming and Formal Methods*, pages 196–212. Springer, 2013.

[13] P. Koopman and M. Wagner. Autonomous vehicle safety: An interdisciplinary challenge. *IEEE Intelligent Transportation Systems Magazine*, 9(1):90–96, 2017.

[14] M. Koschi and M. Althoff. SPOT: A tool for set-based prediction of traffic participants. In *Proc. of the IEEE Intelligent Vehicles Symposium*, pages 1686–1693, 2017.

[15] M. Koschi, C. Pek, M. Beikirch, and M. Althoff. Set-based prediction of pedestrians in urban environments considering formalized traffic rules. In *Proc. of the IEEE Int. Conf. on Intelligent Transportation Systems*, pages 2704–2711, 2018.

[16] S. M. Loos, A. Platzer, and L. Nistor. Adaptive cruise control: Hybrid, distributed, and now formally verified. In *Proc. of the Int. Symposium on Formal Methods*, pages 42–56, 2011.

[17] L. Martinez-Gomez and T. Fraichard. An efficient and generic 2D inevitable collision state-checker. In *Proc. of the IEEE Int. Conf. on Intelligent Robots and Systems*, pages 234–241, 2008.

[18] I. M. Mitchell. Comparing forward and backward reachability as tools for safety analysis. In *Hybrid systems: computation and control*, pages 428–443. Springer, 2007.

[19] C. Pek and M. Althoff. Efficient computation of invariably safe states for motion planning of self-driving vehicles. In *Proc. of the IEEE Int. Conf. on Intelligent Robots and Systems*, pages 3523 – 3530, 2018.

[20] C. Pek and M. Althoff. Computationally efficient fail-safe trajectory planning for self-driving vehicles using convex optimization. In *Proc. of the IEEE Int. Conf. on Intelligent Transportation Systems*, pages 1447–1454, 2018.

[21] C. Pek, M. Koschi, M. Werling, and M. Althoff. Enhancing motion safety by identifying safety-critical passageways. In *Proc. of the 56th IEEE Conference on Decision and Control*, pages 320 – 326, 2017.

[22] C. Pek, M. Koschi, and M. Althoff. An online verification framework for motion planning of self-driving vehicles with safety guarantees. In *AAET - Automatisiertes und vernetztes Fahren*, pages 260–274, 2019.

[23] S. Petti and T. Fraichard. Safe motion planning in dynamic environments. In *Proc. of the IEEE Int. Conf. on Intelligent Robots and Systems*, pages 2210–2215, 2005.

[24] W. Schwarting, J. Alonso-Mora, and D. Rus. Planning and decision-making for autonomous vehicles. *Annual Review of Control, Robotics, and Autonomous Systems*, 1(1):187–210, 2018.

[25] S. Shalev-Shwartz, S. Shammah, and A. Shashua. On a formal model of safe and scalable self-driving cars. *arXiv preprint arXiv:1708.06374*, 1:1–37, 2017.

[26] M. Werling, J. Ziegler, S. Kammel, and S. Thrun. Optimal trajectory generation for dynamic street scenarios in a Frénet Frame. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 987–993, 2010.

[27] J. Ziegler, P. Bender, T. Dang, and C. Stiller. Trajectory planning for Bertha – a local, continuous method. In *Proc. of the IEEE Intelligent Vehicles Symposium*, pages 450–457, 2014.