

General Fail-Safe Emergency Stopping for Highly Automated Vehicles

J. Beyerer^{1B,2}, J. Doll³, F. Duerr⁴, M. Flad^{1C}, M. Frey^{1D}, F. Gauterin^{1D}, S. Hohmann^{1C}, E. Knoch^{1D}, R. Kohlhaas⁵, A. Lauber^{1A}, F. Pistorius^{1A}, M. Roschani², M. Ruf², E. Sax^{1A}, S. Strasser⁶, J. Ziehn^{2,†}

Abstract—From SAE level 3 onwards, automated vehicles must be able to resolve sudden system failures without driver intervention, including failure modes that are difficult or impossible to address by redundancy alone. Causes of hazardous multiple-point faults—beyond internal failures—include lightning strikes or deliberate attacks by electromagnetic pulses. Stopping the vehicle under such conditions is challenging: A full braking maneuver may risk rear-end collisions or loss of traction; likewise, any other constant braking profile will pose considerable risk of *not* achieving a true “safe state”.

This paper presents an emergency stopping system to execute a situation-dependent braking maneuver that can resolve system failures *up to* (but not limited to) a full electronics / electronics failure, with the aim of providing a baseline safety solution for all failure modes (short of mechanical failures) for which no dedicated solution is available.

The system is composed of an electronic planning unit and a hydraulic / mechanical subsystem, both of which are implemented and tested in simulated and in real environments.

I. INTRODUCTION AND STATE OF THE ART

Emergency stopping for automated vehicles comprises a wide range of very different tasks with very different available means of resolution, depending on the nature of the “emergency”. This may include purely external conditions, such as obstacles on the road, as well as inattention or incapacitation of the driver in SAE Level 1 or 2 vehicles, in which the vehicle itself can make full use of all its subsystems but is unable to safely continue its trip as planned (which does not necessarily qualify as an “emergency” by ISO 26262). Examples of approaches to resolve such situations can be found e.g. in [WL12], where model-predictive control is used to avoid collisions with pedestrians; in [MA16] where an emergency maneuver is planned in case of unexpected traffic situations; in [FB11], which provides an overview of emergency collision




Fig. 1: An unoptimized emergency braking can bring the vehicle to a halt on dangerous locations, or risk side or rear-end collisions, loss of traction or passenger injuries. If sufficient free space is available, the vehicle should decelerate as gently as safely possible. The proposed system achieves this behavior for general E/E failures including a *complete electrical failure*.

avoidance algorithms specifically for cooperative vehicles; in [TNK⁺16], where a safety stop system prototype is presented that brings the vehicle to a halt in case of a cardiac emergency of the driver; or in [BCNF14] which describes measures for reaching safe states if the driver fails to intervene on request.

The other extreme is mechanical failures such as damaged wheels or damaged brakes, which are usually not specific to automated driving and very difficult to control for both human drivers and automated systems. They may, however, be less likely for automated vehicles, since wear and tear can be minimized, systematic and predictive maintenance is more easily established, and deviations of nominal conditions can typically be detected earlier (e.g., cf. [MLSD13] for effects of prognostics systems on commercial vehicles).

In between hazardous driving situations and mechanical failures lie emergencies in which the electric and electronic (E/E) subsystems of the vehicle are degraded such that the trip cannot safely be continued, ranging from relatively minor component failures (e.g. in a headlight or in the front radar) up to a potentially unbounded E/E failure, such as by lightning strike or a deliberate attack. Many such failure modes can effectively be mitigated by redundancy [BCNF14], but multiple-point faults remain challenging to resolve, particularly for vehicles for SAE Level 3 and up, which for *controllability* (by ISO 26262: avoiding “harm or damage through the timely reactions of the persons involved”) cannot rely on a “timely reaction” on the part of the driver. A comprehensive overview of approaches to this challenge is given in [Res16].

This paper will focus on the latter type of failures: E/E failure modes that cannot be resolved by any dedicated system or redundancy which require an immediate halt, and which cannot be assumed to be controlled by a human operator.

*This publication was written in the framework of the Profilregion Mobilitätssysteme Karlsruhe, which is funded by the Ministry of Science, Research and the Arts and the Ministry of Economic Affairs, Labour and Housing in Baden-Württemberg and as a national High Performance Center by the Fraunhofer-Gesellschaft.

¹Karlsruhe Institute of Technology (KIT), 76131 Karlsruhe, Germany

^AInstitute for Information Processing Technologies (ITIV)

^BVision and Fusion Laboratory (IES)

^CInstitute of Control Systems (IRS)

^DInstitute of Vehicle System Technology (FAST)

²Fraunhofer IOSB, 76131 Karlsruhe, Germany

³Research Center for Information Technology (FZI), 76131 Karlsruhe, Germany

⁴Audi AG, 85045 Ingolstadt, Germany

⁵Robert Bosch GmbH, Corporate Research, Automated Driving, 71272 Renningen, Germany

⁶TÜV SÜD Auto Partner Ingenieurbüro Mentis, 72070 Tuebingen, Germany

[†]Corresponding author, jens.ziehn@iosb.fraunhofer.de

A. Situation-Dependent Emergency Stopping

A natural and state-of-the-art reaction to the previously described type of failure modes is a full braking maneuver, and reasons for stopping the vehicle as quickly as possible are evident: The lower the speed, duration or distance of any trajectory, the safer it generally is—in particular under the premise of a large-scale system failure. But there are also important reasons for low-dose or delayed emergency braking. Those include clearing dangerous areas (e.g. railroad tracks or intersections), avoiding passenger injuries, avoiding impact of dynamic objects (e.g. in rear-end collisions), and avoiding a state of *dynamic* friction (sliding) which would produce an unstable and unpredictable, potentially even *longer* stopping trajectory (see [KHEL18] for an approach to adapt emergency braking decelerations to the expected road friction).

The choice for an ideal stopping maneuver thus depends on the situation: In low-speed inner-city scenarios on dry roads, a full braking maneuver will typically be safest. On highways with close rear vehicles, or on wet roads, the same action could cause significant accident risks. The means available to achieve the situation-dependent maneuver depend on the worst-case assumption of the emergency system. For example, [RZW⁺15] proposes a solution for stopping the vehicle in a safe location (usually the side of the road) in case of a failed perception-planning pipeline, using only basic electronic control of the actuators and predictions from previous steps.

B. Overview of the Paper

This paper demonstrates the viability of a system for fail-safe emergency stopping that, upon failure, only assumes mechanical and hydraulic system parts to be operational, yet provides a situation-dependent, adaptive stopping maneuver. The aim is to set a lower bound for the performance of a fail-safe emergency stopping that includes a complete electrics / electronics (E/E) failure, and to thereby provide a fallback for a wide range of failure modes where more elaborate safety systems do not apply. Section II gives an overview of the system, and describes and motivates the implementation details of the built prototype. Section III evaluates the prototype both in real test drives on closed-off roads, and in software-in-the-loop (SIL) simulations of traffic scenarios. The results are summarized in Sec. IV.

II. SYSTEM GOALS AND DESCRIPTION


The system is laid out to satisfy the following requirements:

- R1** Bring the vehicle to a full stop in case of an emergency
- R2** Choose a situation-dependent braking maneuver
- R3** Do not use E/E components after failure
- R4** Support integration with existing perception/prediction/planning modules from the regular driving system


To achieve this, the system is composed of two parts, shown in Fig. 3: The planning unit that requires operational electrics and electronics, but is only required *prior* to the system failure; and the hydraulic / mechanical system that assumes control of the car upon failure and brings the vehicle to a halt based on the parameters set by the planning unit. We primarily outline



(a) In a system failure of the blue car, traveling at about 60 km/h (35 mph), a full braking maneuver would risk a rear-end collision with the following car, possibly also a loss of traction, depending on road conditions. If the steering wheel is locked to maintain the current curvature (indicated rainbow line), the vehicle can safely decelerate more gently without entering the highway.



(b) Prediction of collision risks (bright cells: high) over time and arc length of the circular constant-curvature path (extracted from the regular maneuver planning, cf. Fig. 5) and optimization result. The optimization process does not optimize a single maneuver, but a *space* of maneuvers over the unknown exact time of failure until the next planning interval. Depending on when the system fails, different maneuvers would occur. In the example, the system actually fails rather late within the interval, so that the green line occurs.



(c) Result of a physical simulation using the braking pressure ensuing at the simulated failure time indicated in (b) (trajectories of other cars not depicted).

Fig. 2: Example application during highway entry. In (a), the vehicle is still operational; the valve is currently in a state of strong deceleration. The emergency planner determines that for the next interval $[t_{\text{now}}, t_{\text{now}} + \Delta t_{\text{plan}}]$, the valve should change to a gentler deceleration. This results in the subset of trajectories depicted in (b) which occur at different $t_{\text{fail}} \in [t_{\text{now}}, t_{\text{now}} + \Delta t_{\text{plan}}]$. After planning, the vehicle fails, leading to the stopping trajectory in (c).

the design that was eventually implemented and tested in the prototype, but will hint at options for possible alternatives where applicable.

A. Operation Principle

The vehicle is fitted with a hydraulic / mechanical braking system which, upon failure, releases a braking pressure onto




Fig. 3: Overview of the proposed system. Upon failure of the E/E systems (right), detected by a watchdog mechanism **G**, the hydraulic / mechanical subsystem (left) engages. Valve **B** opens and releases the pressure from piston accumulator **A** towards the switch valve **C** that toggles between immediate full-pressure braking (left branch, **E**) or situation-adaptive braking (right branch) dosed by the pressure regulation valve **D**, whose state is adjusted by the emergency planning unit **F** at regular intervals to choose the optimal deceleration profile for the vehicle's current situation.




Fig. 4: Exemplary timing diagram of the developed system. The pressure regulation valve diameter **D** is updated at regular intervals (spaced by Δt_{plan}) by the emergency planning unit **F**. The valve transitions for some time t_{valve} (which depends on the transition step width), during which the system is in an intermediate state between the previous and the next deceleration (hatched areas). When the system fails at some t_{fail} , the watchdog signal **G** ceases, and the lock valve **B** releases the pressure onto the pressure regulation valve **D**, whose state at the moment of failure then determines the braking deceleration.

the wheels that is dosed by an electrically adjustable pressure regulation valve (labeled **D** in Fig. 3). The valve's position α (technically the intended hydraulic pressure, but used here synonymously with an intended braking deceleration for simplicity) is controlled by a planning unit (**F** in Fig. 3).

During operation, the planning unit periodically optimizes the valve position, as shown in Fig. 4. At time t_{now} , a valve position α_{now} is set from the previous step. The optimization process determines an optimal valve position α_{next} to achieve maximum safety if the system fails within the next interval $T_{\text{now}} = [t_{\text{now}}, t_{\text{now}} + \Delta t_{\text{plan}}]$. If the system does not fail within this interval, no braking action is executed; the valve merely moves to α_{next} and the planning process is repeated at $t_{\text{now}} + \Delta t_{\text{plan}}$.

If, on the other hand, the system *does* fail at some $t_{\text{fail}} \in T_{\text{now}}$, detected by a missing watchdog signal, hydraulic pressure is released onto the brakes through the valve, dosed by the

current valve aperture $\alpha(t_{\text{fail}})$. This constant pressure is applied until the vehicle reaches a halt.

Section II-B will describe the design of the planning algorithm that motivates some additional requirements for the hydraulic / mechanical subsystem (beyond the above descriptions), which is subsequently described in Sec. II-C.

B. Planning Unit

The planning unit's main purpose is to find, every Δt_{plan} , an optimal valve transition $\alpha_{\text{now}} \rightarrow \alpha_{\text{next}}$ for the upcoming time interval T_{now} to assure maximum safety if the system fails within T_{now} . To do so, several factors must be considered:

- Since the braking action is engaged if and only if the system fails, the occurrence of the failure can be used as a stochastic condition for the planning; i.e. all steps considering which maneuver to execute may assume that the failure occurs with certainty within T_{now} .

- Beyond this, the time t_{fail} is unknown within T_{now} ; since Δt_{plan} should be very short, prior expectations on when the system is more likely to fail before the next planning step are unlikely. In this case, the probability density $p(t_{\text{fail}})$ can be assumed to be uniform over T_{now} .

- Thus, at t_{now} , no single emergency trajectory can be optimized, but instead a (near) continuous set of trajectories ensuing for different t_{fail} : A later t_{fail} means that the vehicle has traveled further before the failure—and thus the braking—occurs, leading to a stopping position further along the way.

- The shape of this trajectory set is governed by the choice of α_{next} , but valve motion is not instantaneous: The valve will reach its next position not before some $t_{\text{now}} + \Delta t_{\text{valve}}$, and until then, all intermediate states $\alpha(t) \in [\alpha_{\text{now}}, \alpha_{\text{next}}]$ will occur. Realistic values will be $\Delta t_{\text{plan}} \leq 250 \text{ ms}$ and $\Delta t_{\text{valve}} \approx 50 \text{ ms}$, so that the probability of the system failing in an intermediate state is substantial. In this case, the valve stops and the vehicle brakes with $\alpha(t_{\text{fail}})$ between α_{now} and α_{next} .

- The implemented algorithm must be deterministic and real-time capable, such that a solution is provably obtained after a fixed maximum time, namely within Δt_{plan} .

The complete description of the developed planning unit, including a detailed analysis of modeling assumptions, possible extensions and numerical considerations, can be found in [Due18]; we will only provide an overview of the main aspects and considerations for the design of the system.

1) *Connection to the Regular Trajectory Planning System:* To allow integration with existing planning algorithms inside the automated vehicle as shown in Fig. 5—both for re-using predictions and for obeying the same safety criteria—the emergency planning problem is stated in terms of functional optimization, as used for regular maneuver planning e.g. in [ZBDS14], [RZW⁺14]: A *trajectory* is considered a function from time to local planar coordinates $\xi(t) = [x(t), y(t)]^T$, and the planning problem is stated as

$$\xi^* \in \arg \min_{\xi \in \Xi} \mathcal{P}[\xi] \quad \text{with} \quad \mathcal{P}[\xi] = \int_{t_{\text{now}}}^{t_{\text{end}}} dt \ell(\xi(t), \dot{\xi}(t), \ddot{\xi}(t), \dots, t), \quad (1)$$

where ℓ is the *Lagrangian* that assigns penalties, or “costs”, to individual time points along the trajectory, based on po-




Fig. 5: Planning principle, based on the scenario in Fig. 2. The ego vehicle **E** perceives its environment including objects in a static map, and other vehicles **B**. The environment is first used by the regular planner to compute a general prediction for maneuver planning (not shown), to optimize a regular maneuver **C**. The proposed system uses the current steering wheel angle of the ego vehicle to compute a constant-curvature arc **D**, which the vehicle would follow if the steering wheel angle was deliberately locked in an imminent failure. The predicted occupancies and collision risks from the regular maneuver planner are then extracted to yield collision risks **A** (over arc length along **D** and time), including lateral uncertainties as measured in the test drives in Sec. III-A. The two-dimensional risk map **A** is then used to plan the optimal maneuver.

sitions, their derivatives, and time; $\mathcal{P}$ is the *functional* that rates trajectories based on their accumulated penalties over time; and ξ^* , out of a set Ξ of admissible trajectories, is an optimal trajectory (there may be multiple global optima) that minimizes $\mathcal{P}$ and that is to be found by the regular planning algorithm. This formulation allows to integrate goals e.g. for dynamic collision avoidance, comfort / ecology and traffic rules consistently with cost and uncertainty models, and is thus used as the assumed planning input for the emergency planning unit as well. It should also be noted that t_{end} is not $t_{\text{now}} + \Delta t_{\text{plan}}$: Typical applications use a longer prediction horizon (of several seconds) to assure that the maneuver follows a long-term goal. For emergency planning, we may even aim to model *hours* by assigning penalties for unsafe stopping places, such as railroad crossings, even when no train is detected (cf. [Due18, p. 71ff]).

2) *Particularities of Emergency Trajectory Planning*: To integrate the emergency planning problem with the functional optimization statement, we distinguish between the lateral and the longitudinal profile of the trajectory. The first option for a lateral profile is leaving the steering wheel unaffected, to slowly return to zero by mechanical force. This requires no mechanical additions, but gives a complex lateral profile depending i.a. on the force on the front wheels, which in turn depends on the deceleration strength, rendering the planning task laborious. The proposed solution instead assumes that upon failure, the steering wheel is weakly fixed in its current position, such that the lateral profile is circular with good approximation under normal road conditions [RZR⁺14] but the driver is able to override the lock with moderate torque.

We can thus assume that the lateral profile is approximately known during planning; Fig. 3 shows that the regular planning system’s predictions along the constant-curvature path are extracted and passed to the emergency planning unit as arc length over time, as shown in Figs. 2b and 5, allowing for a path-velocity decomposition that simplifies planning [KZ86].

For the longitudinal profile, first note that the emergency trajectory over time depends on the position of the vehicle at t_{fail} and the valve state at t_{fail} , leading to a family $\xi(t, \alpha_{\text{now}}, \alpha_{\text{next}}, t_{\text{fail}})$. Given both α_{now} (which is known), and t_{fail} (which is *not*), we *could* optimize $\mathcal{P}[\xi(t, \alpha_{\text{now}}, \alpha_{\text{next}}, t_{\text{fail}})]$ for α_{next} . Since t_{fail} is uncertain however, we must optimize the expected value $\mathbb{E}\{\mathcal{P}[\xi(t, \alpha_{\text{now}}, \alpha_{\text{next}}, t_{\text{fail}})]\}$ for α_{next} , here using a uniformly random $t_{\text{fail}} \sim \mathcal{U}(t_{\text{now}}, t_{\text{now}} + \Delta t_{\text{plan}})$.

Figure 2b gives an example of the optimization of the continuous set $\xi(t, \alpha_{\text{now}}, \alpha_{\text{next}}, t_{\text{fail}})$ for α_{next} . The optimal *set* shown there includes trajectories that occur if the system fails during valve transition (hatched) as well as trajectories that would occur if the valve reaches α_{next} before failure. Should the system fail at all, any trajectory within the indicated set may occur, so the entire set is placed, solely by choice of α_{next} , such that its containing trajectories optimize expected safety.

Even though only a single parameter α_{next} is optimized, evaluating one particular α_{next} (as the one shown in Fig. 2b) is costly, since a continuum of trajectories must be accumulated. To achieve realtime performance on the low-cost prototype platform, evaluating a dense set of trajectories individually is prohibitive. Instead, the region between the first possible trajectory $\xi(\dots, \alpha_{\text{next}}, t_{\text{fail}} = t_{\text{now}})$ and the last possible trajectory $\xi(\dots, \alpha_{\text{next}}, t_{\text{fail}} = t_{\text{now}} + \Delta t_{\text{plan}})$ is re-parameterized from t_{fail} and α to arc length s , such that the penalties in a re-parameterized Lagrangian $\ell(t, s)$ can be integrated independently of any particular α to obtain $L(t, s)$ with $\partial L / \partial s = \ell$, and the regions swept due to a candidate α_{next} can be evaluated by simple subtractions of the anti-derivative, independently of the number of enclosed trajectories. Care must be taken to consider the different distributions of trajectories with t_{fail} and α , indicated in Fig. 6: While for a constant α , trajectories are indeed spaced (along arc length s) linearly with t_{fail} (Fig. 6b), this does not hold for valve transitions: Since a linear change in acceleration has a quadratic effect on stopping distance, trajectories are spaced non-linearly while the valve is in motion (Fig. 6c). This can lead to the counterintuitive situation that stopping distances can *shorten* for later t_{fail} , if the valve in this time moves to stronger decelerations (Fig. 6d). An exhaustive analysis of the issue, and steps to achieving around 30 ms planning time on a lightweight computer platform (Raspberry Pi 3B, using a single ARM Cortex A53 core at 1.2 GHz), can be found in [Due18].

3) *Masking Hazardous States*: The previously described system has a significant limitation that can be illustrated on the example of a railroad crossing as seen in Fig. 1: If the vehicle approaches the railroad crossing, the planning algorithm would first aim to stop the vehicle *before* the crossing; at some point, when the vehicle cannot safely be stopped before reaching the tracks, the planner will switch to




Fig. 6: Key effects on the stopping trajectory to be considered in the planning stage. The basic system goal is to vary the deceleration to adjust the vehicle’s stopping position (a). However, due to the unknown moment of failure over the planning interval (b) and the finite speed of the valve, which traverses all decelerations from the current deceleration to the next over some time (e.g. 40 ms in c–d), the actual effects of changing to a different deceleration are complex and possibly counterintuitive. The distribution of endpoints also hints at the density distribution of trajectories that must be considered during planning in Sec. II-B, where no single trajectory can be planned, but instead the uncertain t_{fail} gives rise to a family of possible trajectories.

a new valve state that stops the vehicle *behind* the tracks. If the vehicle fails during this transition, the system will bring it to a halt at the very position it should avoid. This points to a general issue: Due to the risk of failure during valve transition, transition intervals (the hatched areas) cannot safely cross dangerous stopping positions; the chances of stopping there are significantly reduced, but not eliminated. To address this issue, a second valve is added: A switch valve (C in Fig. 3) can toggle between a constant braking deceleration (e.g. a full braking) and the function of the pressure regulation valve (D), to *mask* the valve’s transition—if necessary. The planning algorithm is extended to identify such inevitable transitions using a dynamic programming approach. If it finds the transition interval to be safe (e.g. when the car is initially approaching the train tracks from a long distance), the output from the pressure regulation valve remains connected to the braking cylinder. If an unsafe transition is unavoidable, the switch valve changes to the constant braking default while the pressure regulation valve moves, and redirects to the pressure regulation valve only when the valve has reached α_{next} . Since the switch valve can be chosen to be fail-open, it cannot fail in transition. The extension, again laid out in more detail in [Due18], requires additional computational effort, raising execution times to around 120 ms on the Raspberry Pi, which still enables planning rates of 8 Hz.

C. Hydraulic / Mechanical Subsystem

The hydraulic / mechanical subsystem, partially described in [Wei17], brings the vehicle to a halt based on the parameters determined during planning (Sec. II-B). The pressure regulation valve (D in Fig. 3), whose position is set by the planning system, is connected to a piston accumulator (A) via two switch valves (B and C). Switch valve B is fail-open and connected to the watchdog system, to release the pressure in the piston accumulator upon failure. Switch valve C realizes the “masking” feature outlined in Sec. II-B3, where hazardous

transitions of the pressure regulation valve are skipped by switching to a constant braking default via path E.

The pressure regulation valve D is fail-stationary, such that upon system failure, when the pressure from A is released, the previous aperture diameter persists. Contrary to this, valve B is fail-open towards A, and switch valve C is fail-open towards B, such that B initiates braking as soon as the system fails, and C maintains the constant braking setting unless the system fails strictly after a new safe α_{next} has been reached.

III. TESTING

The system was evaluated under two aspects: The physical performance of the hydraulic / mechanical system introduced in Sec. II-C, which was not tested on public roads but on closed-off sites of the Test Area Autonomous Driving Baden-Württemberg¹ [FDW⁺19] (Sec. III-A), and the functional performance of the planning algorithm introduced in Sec. II-B, tested in the loop with the OCTANE² platform, using simulated traffic scenarios, with the physical behavior modeled according to the conducted test drives (Sec. III-B).

A. Physical Tests

The main objective of physical testing is determining the accuracy to which a vehicle’s path can be predicted ahead of the failure, given that under the assumed failure mode, no active, closed-loop control can be used to correct the vehicle’s trajectory afterwards.


While uncertainties of any order can explicitly be included in the planning algorithm, exceeding levels of uncertainty imply little gain from the system compared to a classical full braking default. For practical reasons, the physical tests of the longitudinal and lateral accuracy were conducted in three different setups and in two different test vehicles: The lateral accuracies were tested in an electrical VW e-Golf 7 *without* the

¹www.taf-bw.de


²www.octane.org



(a) Experiments for longitudinal accuracy with prototype installed in the VW Golf 7 Variant test vehicle. (b) Hydraulic / mechanical prototype in the back of the vehicle.



(c) Results for longitudinal accuracy with the prototype by stopping from three velocities (black: ≈ 30 km/h, blue: ≈ 50 km/h, light green: ≈ 70 km/h) and using three different valve parameters, aligned at stopping position and time. Accuracies are given as distance intervals. Each test includes five or more trajectories except for the medium deceleration from 70 km/h (marked with an asterisk) which only has three measurements.



(d) Experiments for lateral accuracy conducted with a VW e-Golf 7 on a non-public site of the Test Area Autonomous Driving BW. (e) Experimental predictability of a vehicle's lateral position at constant steering angle $\delta = -180^\circ$ with v_0 up to 30 km/h and a down to -5 m/s 2 .

Fig. 7: Physical test drives for longitudinal (a–c) and lateral accuracy (d–e).

prototype installed, using by-wire control and manual control, at different sites for determining the effects of high curvatures (Sec. III-A1) and long stopping distances (Sec. III-A2). The longitudinal test drives (Sec. III-A3) were conducted with the prototype installed into a VW Golf 7 Variant. Both test vehicles have a near-identical setup [FR17] beyond the stated differences, with trajectory measurements using RTK GPS at centimeter accuracy, along with onboard odometry and IMUs.

1) *Lateral High Curvature Tests:* The test drives to establish the lateral accuracy to which a vehicle's path can be predicted at constant high steering angles were conducted at KIT Campus East. The steering wheel angle was kept at a constant 180° while the vehicle was operated at speeds of up to 30 km/h and decelerations down to $a = -5$ m/s 2 . At a turn radius of about 14 m, the initial conditions already include strong lateral accelerations of 5 m/s 2 , which are unlikely to occur in typical automated vehicle operation. The results of $n = 50$ test trajectories, shown to scale with lateral standard deviations in Fig. 7e, indicate that even under such dynamics, predictability is reasonable with accuracies strictly below 2 m.

2) *Lateral Long Distance Tests:* The test drives to establish the lateral accuracy at long stopping distances were conducted at the Karlsruhe fairgrounds, which has an uneven surface with sawtooth drainage gradients of approximately 20 m period and 1% slope. Here the vehicle was accelerated to speeds between 10 km/h and 80 km/h and then stopped using decelerations between -2 m/s 2 and -10 m/s 2 . The steering wheel angle was manually held at neutral, with recorded data showing variations within $\pm 3^\circ$. The resulting accuracies of $n = 36$ test drives, again to scale, are given in Fig. 8, with typical German highway lane widths of 3.5 m [For08] indicated. Lateral standard deviations of less than 1 m at stopping distances of 90 m suggest workable levels of uncertainty, with the vehicle entirely inside its lane (w.r.t. the standard deviation) up to a stopping distance of 80 m.

3) *Longitudinal Tests:* The hydraulic / mechanical prototype, shown in Fig. 7b, was installed in a VW Golf 7 Variant and tested on a long straightaway at KIT Campus North, Karlsruhe. Due to the system assumption of an unbounded E/E failure, ABS and ESC were hard disabled, as these would be unavailable under actual failure conditions either. This entailed a lack of speed readings from the CAN, so that the car was manually accelerated to approximate GPS speeds of 30 km/h, 50 km/h and 70 km/h, and decelerated using three fixed valve parameters α . Results of $n = 70$ trajectories are shown in Fig. 7c; for better comparability, all obtained trajectories are aligned by their endpoints and stopping times, and accuracy




Fig. 8: Lateral accuracy test drives over long distances. Standard deviations are indicated along with a regular German highway lane width for scale.

intervals are indicated at the earliest common start points. On the given favorable road surface conditions, longitudinal accuracies of better than 6m are obtained, which suggest that the prototype itself is capable of achieving reproducible braking that can be used in stochastic predictions.

B. Simulation Tests

The simulation scenarios (more of which can be found discussed in more detail in [Due18]) are set up to test the functionality of the planning unit. We show two scenarios: Failure during highway entry and failure during an urban left-turn maneuver. Different failure times throughout each scenario were analyzed, with the exact instant of failure, t_{fail} , being drawn randomly after planning.

The prediction system outlined in [RZW⁺14] is used as the regular maneuver planning system, providing the collision risks over arc length and time for the emergency planner as shown in Fig. 5. The planning system is tested in a continuous loop on the Raspberry Pi platform for the entire duration of each scenario (i.e. continuously over several planning cycles), such that initial valve states in each planning cycle are the actual results of previous planning cycles. For example, the valve in Fig. 2b is initially in a stronger deceleration position because in the previous step, the steering wheel angle was near-neutral and a long stopping distance would have caused a lane departure.

After each planning step, a random time $t_{\text{fail}} \in T_{\text{now}}$ is drawn, and the failure, including the resulting emergency trajectory, is then simulated. The trajectory accuracies, and the physical modeling of the braking maneuver, are calibrated to the test drives with the VW e-Golf 7 from Sec. III-A.

Some simulation results are shown in Figs. 9. Each simulation scenario shows the initial situation, the predictions from the regular planning system, and the planned emergency maneuver set. The subsequently simulated event of failure is plotted as a green line ending in a triangle, indicating the *expected* failure maneuver at this t_{fail} , and as a rendering of the resulting maneuver obtained by a physical simulation of the vehicle motion with the resulting braking pressure.

Quantitative safety metrics would depend significantly on realistic behavior and reaction models for other traffic participants, which are currently unavailable; yet, the selected scenarios indicate that the planning models are capable of achieving solutions that increase safety compared to any constant braking deceleration, by maximizing braking distances (and thus e.g. tire friction and reaction times of rear vehicles) wherever possible, while successfully avoiding unsafe areas.

IV. CONCLUSION AND OUTLOOK


We have presented an emergency stopping system for highly automated vehicles that is able to resolve failure modes up to a complete electronics / electronics (E/E) failure, while providing a situation-dependent and predictive braking maneuver. The system is composed of a hydraulic / mechanical branch, designed to provide an accurate fail-safe stopping maneuver for any moment of failure, and a stochastic planning algorithm that controls the hydraulic / mechanical parameters prior to




(a) In this continuation of the scenario in Fig. 2, the ego vehicle stayed operational until entering the highway. Once it has reached zero steering angle, it can pick a gentle deceleration to minimize the risk of a rear-end collision. In the simulated failure, the vehicle fails while the valve is in motion (the hatched area). Since the transition to a gentler deceleration here continuously increased safety, it was not masked as described in Sec. II-B3.



(b) In this inner city scenario (continued in the following frames), the ego vehicle is followed closely by another car. With a free straight path, a gentle deceleration minimizes the risk of a rear-end collision as above.



(c) Once the vehicle starts steering left, the gentle deceleration **A** risks a head-on collision on the opposite lane. The system switches to immediate full braking **B** while moving the valve to an intermediate range **C** (cf. Sec. II-B3), where the system actually fails. The vehicle keeps clear of the opposite lane.



(d) One step later the planner finds a trajectory across the two opposite lanes to a safe stopping location. To safely move the valve from the current deceleration **A** to the goal **C**, it again first switches to a full braking, to avoid failing at an intermediate deceleration between **A** and **C** that would place it directly in the path of the oncoming vehicles. The system actually fails rather early within T_{now} , resulting in a relatively short but safe braking distance.

Fig. 9: Simulation results each showing the initial situation, the planning (over a prediction horizon of 10 s each), and a physically simulated failure trajectory.

failure, requires little additional computational effort, and can be connected to regular maneuver planners to re-use prediction results and planning criteria.

The design parameters of both the algorithmic and the hydraulic / mechanical parts were outlined and motivated. A prototype of the system was built and installed in a test vehicle, and tested separately in SIL simulations for the functional validity of the planning algorithm, and on closed test areas for the physical performance of the hydraulic / mechanical system. The results, while clearly not comprehensive enough to conclusively determine quantitative system parameters, indicate that the system can execute adaptive stopping maneuvers with good accuracy, and that the planning algorithm provides adequate solutions even in challenging scenarios.

Outlook

Several enhancements for the system were considered but not yet integrated into the prototype. A physical lock that fixes the steering wheel angle upon failure, while still allowing manual override, was not yet built. If integrated, the planning system could be extended to decide whether to lock the wheel upon failure, or whether to leave it to return to neutral by itself, providing another degree of freedom.

Furthermore, extensions to the planning module are conceivable that improve result quality, or computational effort. If two pressure regulation valves were used instead of one, and a switch valve between them, it would be possible to modify the maneuver near-continuously by alternating between the valves and always moving the disconnected one. This would eliminate the risk of failing while the valve is in motion. If the pressure release upon failure was physically blocked until the end of the planning interval T_{now} , the moment of braking could be determined with certainty.

A combination of both means increases the complexity of the hydraulic / mechanical system (and was therefore not pursued), but makes the planning task easier since uncertainty is reduced considerably, which may also produce safer results. Besides varying the braking pressure, other variable parameters could be considered to affect the trajectory, including a variable delay before the onset of braking, or a time-variable braking profile.

As previously noted, the evaluation of the system primarily serves as a proof of concept. According to the motivation, it is assumed that gentler decelerations compare favorably to a full braking default both in adverse road conditions and with distracted other traffic participants; yet, a more thorough evaluation on the predictability of the stopping trajectories for various driving and road surface conditions is required, as well as simulation tests using realistic behavior of other traffic participants, to establish a quantitative estimate of the safety gain provided by the system.

REFERENCES

- [BCNF14] J. Becker, M.-B. A. Colas, S. Nordbruch, and M. Fausten. Bosch's vision and roadmap toward fully autonomous driving. In *Road vehicle automation*, pages 49–59. Springer, 2014.
- [Due18] F. Duerr. Notbremssystem für Systemausfälle im vollautomatischen Fahren. Master's thesis, Karlsruhe Institute of Technology KIT, Fraunhofer IOSB, Karlsruhe, Germany, April 2018.
- [FB11] Ch. Frese and J. Beyerer. A comparison of motion planning algorithms for cooperative collision avoidance of multiple cognitive automobiles. In *2011 IEEE Intelligent Vehicles Symposium (IV)*, pages 1156–1162. IEEE, 2011.
- [FDW⁺19] T. Fleck, K. Daaboul, M. Weber, P. Schörner, M. Wehmer, J. Doll, S. Orf, N. Sußmann, Ch. Hubschneider, M. R. Zofka, F. Kuhnt, R. Kohlhaas, I. Baumgart, R. Zöllner, and J. M. Zöllner. Towards large scale urban traffic reference data: Smart infrastructure in the test area autonomous driving baden-württemberg. In M. Strand, R. Dillmann, E. Menegatti, and S. Ghidoni, editors, *Intelligent Autonomous Systems 15*, pages 964–982, Cham, 2019. Springer International Publishing.
- [For08] Forschungsgesellschaft für Straßen- und Verkehrswesen e. V. (FGSV). Richtlinien für die Anlage von Autobahnen (RAA). Technical report, FGSV, 2008.
- [FR17] M. Filsinger and M. Roschani. Strategisches Invest: Versuchsfahrzeuge für Technologie-Experimente (VERTEX). *visIT Autonome Mobilität*, pages 12–13, January 2017. <https://www.iosb.fraunhofer.de/servlet/is/69530/>.
- [KHEL18] I. Koglbauer, J. Holzinger, A. Eichberger, and C. Lex. Autonomous emergency braking systems adapted to snowy road conditions improve drivers' perceived safety and trust. *Traffic injury prevention*, 19(3):332–337, 2018.
- [KZ86] Kamal Kant and Steven W Zucker. Toward efficient trajectory planning: The path-velocity decomposition. *The international journal of robotics research*, 5(3):72–89, 1986.
- [MA16] S. Magdici and M. Althoff. Fail-safe motion planning of autonomous vehicles. In *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, pages 452–458. IEEE, 2016.
- [MLSD13] M. Mesgarpour, D. Landa-Silva, and I. Dickinson. Overview of telematics-based prognostics and health management systems for commercial vehicles. In *International conference on transport systems telematics*, pages 123–130. Springer, 2013.
- [Res16] A. Reschka. Safety concept for autonomous vehicles. In *Autonomous Driving*, pages 473–496. Springer, 2016.
- [RZR⁺14] M. Ruf, J. Ziehn, B. Rosenhahn, J. Beyerer, D. Willersinn, and H. Gotzig. Evaluation of an Analytic Model for Car Dynamics. In *International Conference on Mechatronics and Control (ICMC), Jinzhou, China*, pages 2446–2451, July 2014.
- [RZW⁺14] M. Ruf, J. Ziehn, D. Willersinn, B. Rosenhahn, J. Beyerer, and H. Gotzig. A Continuous Approach to Autonomous Driving. In *Vehicle and Infrastructure Safety Improvement in Adverse Conditions and Night Driving (VISION), Versailles*, October 2014.
- [RZW⁺15] M. Ruf, J.R. Ziehn, D. Willersinn, B. Rosenhahn, J. Beyerer, and H. Gotzig. Global Trajectory Optimization on Multilane Roads. In *18th IEEE International Conference on Intelligent Transportation Systems (ITSC)*, September 2015.
- [TNK⁺16] I. Takahashi, T.T. Nguyen, H. Kanamori, T. Tanaka, Sh. Kato, Y. Ninomiya, E. Takeuchi, T. Nakagawa, M. Makiguchi, and H. Aoki. Automated safety vehicle stop system for cardiac emergencies. In *2016 IEEE International Conference on Emerging Technologies and Innovative Business Practices for the Transformation of Societies (EmergiTech)*, pages 9–12. IEEE, 2016.
- [Wei17] S. Weissenbach. Konzeption und Aufbau eines Funknothaltesystems für autonom fahrende Elektrofahrzeuge. Master's thesis, Karlsruhe Institute of Technology KIT, Karlsruhe, Germany, September 2017.
- [WL12] M. Werling and D. Liccardo. Automatic collision avoidance using model-predictive online optimization. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 6309–6314. IEEE, 2012.
- [ZBDS14] J. Ziegler, P. Bender, T. Dang, and Ch. Stiller. Trajectory planning for Bertha – A local, continuous method. In *Proceedings of the 2014 IEEE Intelligent Vehicles Symposium (IV), Dearborn*, pages 450–457, June 2014.