

Contingency Model Predictive Control for Automated Vehicles

John P. Alsterda^{1,2}, Matthew Brown¹ and J. Christian Gerdes¹

Abstract— We present Contingency Model Predictive Control (CMPC), a novel and implementable control framework which tracks a desired path while simultaneously maintaining a contingency plan – an alternate trajectory to avert an identified potential emergency. In this way, CMPC anticipates events that *might* take place, instead of reacting when emergencies occur. We accomplish this by adding an additional prediction horizon in parallel to the classical receding MPC horizon. The contingency horizon is constrained to maintain a feasible avoidance solution; as such, CMPC is *selectively* robust to this emergency while tracking the desired path as closely as possible. After defining the framework mathematically, we demonstrate its effectiveness experimentally by comparing its performance to a state-of-the-art deterministic MPC. The controllers drive an automated research platform through a left-hand turn which may be covered by ice. Contingency MPC prepares for the potential loss of friction by purposefully and intuitively deviating from the prescribed path to approach the turn more conservatively; this deviation significantly mitigates the consequence of encountering ice.

I. INTRODUCTION

In a future when automated vehicles are ubiquitous, we will expect them to handle the surprises and challenges that humans tackle routinely. They must avoid children who may run into the street, or maneuver safely when encountering an icy surface. In control engineering, this concept is known as robustness: a controller’s ability to safely operate when subject to unexpected variation in system or environmental parameters. In the field of Model Predictive Control (MPC), work to this end falls largely into two broad categories: Robust MPC (RMPC) and Stochastic MPC (SMPC).

In two reviews, Mayne described some RMPC strategies, their limitations, and direction for future development [1][2]. In these robust algorithms, uncertain parameters’ values are first confined to a set between upper and lower bounds. A control trajectory is then calculated which satisfies system constraints for *all* possible combinations and realizations of these parameters. Thus RMPC is sometimes dubbed ‘worst case’ MPC, prepared for the most extreme coincidences of bad luck – however unlikely.

RMPC is often cast as a min-max problem, solving for commands that minimize the maximum possible cost [3]. The problem is in general nonlinear and non-convex, but Sartipizadeh *et al.* applied this technique on an approximate convex hull of uncertainty to achieve computational feasibility [4]. D. He *et al.* also ensured the tractability of their RMPC work by employing ‘quasi-min-max’ MPC [5].



Fig. 1. Volkswagen GTI Automated Test Platform

Another RMPC instantiation is Tube MPC, which reduces pessimism by solving for control policies instead of open-loop control sequences [1]. Closed-loop policies can react to future events and therefore pursue higher performance in the near term. Tube MPC attempts to guarantee the set of possible disturbed states is contained within problem constraints [6]. Kögel *et al.* presented a formulation to further reduce conservatism by combining the effects of estimation error and model uncertainty [7].

Robust MPC remains a principled strategy for systems that require *guaranteed* stability, like a nuclear reactor [8]. However, it can be impractically cautious when uncertainties are numerous and large, as is the case driving on legal roads. In this article, a *selectively* robust model predictive controller is introduced which achieves responsible but practical conservatism in two ways: First, contingency scenarios package uncertainties that are likely to occur together. Whereas traditional RMPC assumes all uncertainties may agitate simultaneously, engineers can leverage expert knowledge or parameter covariance to focus on the most likely emergency events. Second, we separate nominal and contingency planning into separate horizons, and define a unique cost function for each. By separating nominal and emergency costs, the contingency trajectory is unburdened by objectives typically found in MPC cost functions. Whereas traditional RMPC would aim to maintain comfortable and efficient emergency maneuvers, Contingency MPC preserves only a feasible avoidance trajectory.

Stochastic MPC is another MPC technique to handle uncertainty. Control engineers seeking robustness with less pessimism favor building parameter ambiguity and stochastic uncertainty directly in their optimizations [9]. In SMPC, uncertain quantities are characterized into statistical distributions, and control inputs are calculated to minimize *expected* cost. In this way, the stochastic algorithms do not place undue focus on highly improbable scenarios. Farina *et al.*

¹Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA

²Corresponding author: alsterda@stanford.edu

characterized the general problem statement, which requires a disturbance model to be defined either by a probability density function (e.g. the Gaussian) or some function which produces samples [10]. This assumption, however, may not hold for some systems or applications.

For example, a well-defined distribution may not always exist for the coefficient of friction along a roadway. We may not feel comfortable assigning a probability density to predict a child's behavior. On the contrary, if ice is likely to exist somewhere on a winter day, we should drive in such a way that a surprise icy patch can be handled safely. If children are seen playing on the sidewalk, we should anticipate the possibility they run into the street and approach with an avoidance strategy – a contingency plan.

One SMPC technique has been borrowed to formulate Contingency MPC. In Scenario SMPC, a prediction horizon tree is built which branches from stage-nodes by sampling from a parameter distribution. At each stage-node only one command is calculated, an important constraint to prevent Scenario SMPC from improper prognostication. Lucia *et al.* demonstrated this technique with a nonlinear model, and Krishnamoorthy *et al.* worked to improve the distributed systems over which Scenario SMPC can be decomposed [11] [12]. Contingency MPC's prediction horizon structure resembles a simple scenario tree, with a single stage-node at x_0 . At this stage, we also employ the single-command constraint, shown in Fig. 2. This constraint is central to the function of CMPC; the input u^0 is calculated to be valid for both nominal and contingency horizons. Thus the vehicle will track the desired path as closely as possible, while maintaining the ability to choose the contingency trajectory.

This article offers a new option for control engineers who seek robustness from foreseeable but uncertain events. In Section II, the general Contingency MPC problem is defined. Section III narrows the general problem to an linear convex program for use in automated vehicle (AV) control. In Section IV, we present an experimental demonstration of CMPC. First the test platform turns through a snow-covered corner; it then encounters ice performing the same maneuver. Finally, Section V explores the potential for CMPC to tackle other problems in control under uncertainty and outlines future steps to develop the framework.

II. CMPC FORMULATION

Model Predictive Control is a receding horizon optimal control technique whose central strength is its ability to minimize future cost while explicitly respecting model and environmental constraints [13]. At each time step, MPC calculates a prediction horizon (open-loop trajectory) of discrete states x and inputs u to optimize a cost function, subject to constraints. Upon computation, u^0 is deployed and a new horizon is optimized. Contingency MPC's receding horizons are illustrated in Fig. 2, where x^0 represents the plant's current state and two trajectories are calculated. Trajectory x_{nom} is analogous to the traditional deterministic MPC horizon, containing the states we intend to guide the vehicle through by application of control inputs u_{nom} . On

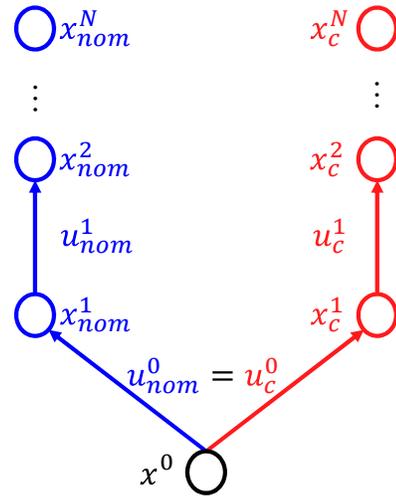


Fig. 2. Contingency MPC prediction horizon. A nominal and emergency trajectory are computed, which share a zeroth action.

the other hand, trajectory x_c represents our contingency plan of action, subject to a unique set of constraints and costs.

A. General Case

The following optimization problem calculates N control inputs $\mathbf{u} \in \mathbb{R}^{2 \cdot m}$ to guide state vectors $\mathbf{x} \in \mathbb{R}^{2 \cdot n}$ to minimize cost function J while abiding to equality constraints H and inequality constraints G . N is the number of prediction horizon stages, n the length of the nominal state vector, and m the number of control inputs.

$$J(x^0) = \min_{\mathbf{u}} \sum_{k=0}^N j^k(\mathbf{x}^k, \mathbf{u}^k) \quad (1a)$$

where:

$$\mathbf{x}^k = \begin{bmatrix} x_{nom} \\ x_c \end{bmatrix}^k ; \quad \mathbf{u}^k = \begin{bmatrix} u_{nom} \\ u_c \end{bmatrix}^k \quad (1b)$$

such that:

$$h(\mathbf{x}, \mathbf{u}) = 0 \quad \forall \quad h \in H \quad (1c)$$

$$g(\mathbf{x}, \mathbf{u}) \leq 0 \quad \forall \quad g \in G \quad (1d)$$

$$\text{and for } k = 0: \quad u_{nom}^0 = u_c^0 = u^0 \quad (1e)$$

Equality constraints H include the dynamics model:

$$\mathbf{x}^{k+1} = f(\mathbf{x}^k, \mathbf{u}^k) \quad \forall \quad k \quad (1f)$$

As in typical MPC formulations, the cost function terms j can be used to encourage the nominal states x_{nom} to follow a desired trajectory, maintain comfort or smoothness, minimize fuel use, etc. However, these costs need not be levied against the contingency trajectory. The separation of costs allows the contingency trajectory to focus on maintaining a feasible emergency maneuver, unburdened by costs unimportant during crisis. Constraint sets H and G encode the contingency scenario. For example, constraints on x_c can be added to represent an obstacle which may enter the road. Alternatively,

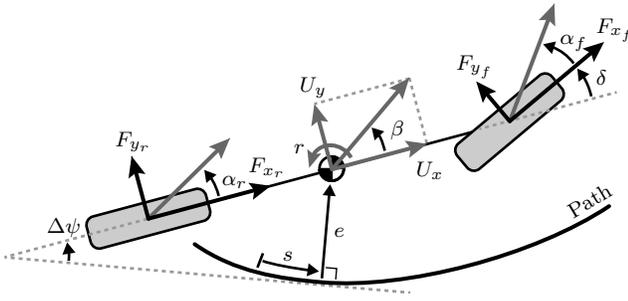


Fig. 3. Bicycle model diagram

to prepare for a dynamics-based contingency, the propagation model for x_c^{k+1} can be modified from the nominal case.

The resulting optimized horizon contains two input trajectories from x^0 : u_{nom} and u_c . As in classical MPC, u^0 is deployed – the root command for both u_{nom} and u_c . Therefore the vehicle *never chooses* between the nominal and contingency trajectories. u_0 tracks the desired path to the greatest extent possible, while maintaining an avoidance maneuver for the contingent event. If the contingency event occurs, CMPC is expected to return safe trajectories in subsequent time-steps. Alternatively if the contingency ceases to be relevant, the vehicle can proceed along the desired path.

III. LINEAR CMPC FOR AV CONTROL

The general CMPC formulation above can be adapted to any MPC framework. For the remainder of this article, the general form is narrowed to a convex linear optimization extended from the deterministic MPC controller developed by Brown *et al.*, and modified to solve directly for steering angle δ as presented by Zhang *et al.* [14] [15].

A. Vehicle Dynamics Model

To adequately handle contingency scenarios which push the limits of vehicle handling, dynamics must be modeled with an appropriate degree of fidelity. We accomplish this using a planar bicycle model with successively linearized tire forces to represent the vehicle [16]. Fig. 3 illustrates this model with three position and velocity states, and forces at each axle. The position states are local to a path with curvature κ , where s represents the vehicle's longitudinal progress, e its lateral error from the path, and $\Delta\psi$ is the heading angle error. Velocity states are composed of longitudinal speed U_x , lateral speed U_y , and yaw rate r . The following differential equations govern the states' evolution:

$$\dot{s} = U_x - U_y \Delta\psi \quad (2a)$$

$$\dot{e} = U_y + U_x \Delta\psi \quad (2b)$$

$$\dot{\Delta\psi} = r - \kappa U_x \quad (2c)$$

$$\dot{U}_x = \frac{F_{xf} + F_{xr}}{m} + r U_y \quad (2d)$$

$$\dot{U}_y = \frac{F_{yf} + F_{yr}}{m} - r U_x \quad (2e)$$

$$\dot{r} = \frac{aF_{yf} - bF_{yr}}{I_z} \quad (2f)$$

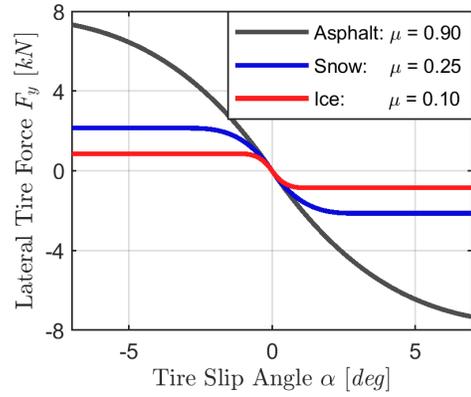


Fig. 4. Fiala brush tire models for asphalt, snow, and ice conditions

Forces F_{yf} and F_{yr} are modeled by a Fiala brush tire curve which relates lateral force to tire slip angles [17]:

$$\delta + \alpha_f = \tan^{-1} \left(\frac{U_y + ar}{U_x} \right) \quad (3a)$$

$$\alpha_r = \tan^{-1} \left(\frac{U_y - br}{U_x} \right) \quad (3b)$$

In Fig. 4, tire models are plotted for asphalt, snow, and ice conditions, which are the coefficients μ used in the experiments which follow.

As a first demonstration of CMPC, this formulation calculates only the steering angle δ . Throttle and braking commands are calculated upstream by a feedforward-feedback scheme, as developed by Funke *et al.* [18]. Including longitudinal forces into CMPC is an opportunity for future development. Thus the CMPC problem we formulate sees F_{xf} , F_{xr} , U_x , and s as a known profile along the prediction horizon. The MPC state vector is then $x = [U_y \ r \ \Delta\psi \ e]^T$.

To complete the MPC dynamics model, the equations of motion and tire model are linearized with respect to operating points from the previous iteration's solution. The linear model was then discretized with respect to δ : A zero-order hold was used for the first ten 20 ms steps, and a first-order hold was used for forty 300 ms steps to form a twelve second horizon. The number and length of time steps were chosen to capture high frequency dynamics in the short term, while extending the total horizon such that CMPC can *see* far enough to prepare avoidance maneuvers. These choices were also influenced by computation time, as reported in Section IV. After discretization, the resulting affine dynamics equations take the form [15]:

$$x^{k+1} = A^k x^k + B^k \delta^k + C^k \quad (4a)$$

To support both nominal and contingency prediction horizons, (4a) is duplicated for both x_{nom} and x_c state vectors:

$$\begin{aligned} \mathbf{x}^{k+1} &= \begin{bmatrix} x_{nom} \\ x_c \end{bmatrix}^{k+1} \\ &= \begin{bmatrix} A_{nom} & 0 \\ 0 & A_c \end{bmatrix}^k \mathbf{x}^k + \begin{bmatrix} B_{nom} & 0 \\ 0 & B_c \end{bmatrix}^k \mathbf{u}^k + \begin{bmatrix} C_{nom} \\ C_c \end{bmatrix}^k \end{aligned} \quad (4b)$$

B. Linear CMPC Problem Statement

The following optimization is extended from the deterministic MPC presented in [14] to calculate a smooth trajectory which follows a desired path while adhering to dynamics, stability, and environmental constraints.

$$\min_{\mathbf{u}} (\mathbf{x}^{50})^T \begin{bmatrix} 0 & 0 \\ 0 & Q \end{bmatrix} \mathbf{x}^{50} + \sum_{k=0}^{50} (\mathbf{x}^k)^T \begin{bmatrix} Q & 0 \\ 0 & 0 \end{bmatrix} \mathbf{x}^k + (\mathbf{v}^k)^T \begin{bmatrix} R & 0 \\ 0 & 0 \end{bmatrix} \mathbf{v}^k + W \begin{bmatrix} \sigma_{stab}^k \\ \sigma_{env}^k \end{bmatrix} \quad (5a)$$

where:

$$Q = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} ; \quad R = 0.01 \quad (5b)$$

$$\mathbf{v}^k = \begin{bmatrix} (u^k - u^{k-1})_{nom} \\ (u^k - u^{k-1})_c \end{bmatrix} ; \quad W = \begin{bmatrix} 50 & 0 \\ 0 & 500 \end{bmatrix}$$

Q applies cost to heading and lateral state errors, R penalizes change in steering angle (slew rate), and W discourages stability and environmental constraint violation. Weights were tuned for experimental performance and chosen to establish priority among the problem's goals: Most important is driving the constraints' slack variables to zero to ensure stability and obstacle avoidance. Next, the nominal trajectory should adhere to the desired path. Lastly the nominal input trajectory should be smooth and comfortable for passengers. This minimization is subject to (4b) and the following constraints. First the steering angle and its slew rate are bounded:

$$|\mathbf{u}^k| \leq \begin{bmatrix} \delta_{max} \\ \delta_{max} \end{bmatrix} ; \quad |\mathbf{v}^k| \leq \begin{bmatrix} v_{max} \\ v_{max} \end{bmatrix} \quad \forall k \quad (5c)$$

Next, yaw rate and sideslip are confined to an invariant set, as developed by Bobier *et al.* and Beal *et al.* [19] [20]. This stability envelope is illustrated in Fig. 7 and enforces $|r| \leq \mu g / U_x$ and $|\beta| \leq \alpha_{rear,peak} + br / U_x$:

$$\begin{bmatrix} H_{nom} & 0 \\ 0 & H_c \end{bmatrix}_{stab} \mathbf{x}^k \leq \begin{bmatrix} G_{nom} \\ G_c \end{bmatrix}_{stab} + \sigma_{stab} \quad \forall k \quad (5d)$$

To prevent collision with road edges, an environmental envelope encodes the inequality $e_{min} \leq e \leq e_{max}$:

$$\begin{bmatrix} H_{nom} & 0 \\ 0 & H_c \end{bmatrix}_{env} \mathbf{x}^k \leq \begin{bmatrix} G_{nom} \\ G_c \end{bmatrix}_{env} + \sigma_{env} \quad \forall k \quad (5e)$$

Finally the nominal and contingency trajectories are coupled via their first commands:

$$u_{nom}^0 = u_c^0 \quad (5f)$$

In the cost function and inequality constraints, block diagonal matrices are employed to weight and constrain the nominal and contingency trajectories independently. For

example, the terminal cost on \mathbf{x}^{50} applies penalty only to the contingency trajectory's final angular and lateral deviation, an important feature to mitigate the shortcomings of a finite horizon. No path following or input-smoothing penalty is applied to the contingency trajectory for $k < N$, however, because we only require it to obey the imposed constraints.

IV. EXPERIMENTAL RESULTS AND ANALYSIS

Contingency MPC was demonstrated in experiment using the convex linear formulation presented in Section III, deployed on an automated Volkswagen GTI (Fig. 1). These results compare CMPC with a deterministic MPC based on the nominal costs and constraints. In each comparison, CMPC purposefully deviated from the desired path to maintain an avoidance strategy for the anticipated contingency.

The first setting is shown in Fig. 5, a snow covered left-hand turn driven at 5 m/s – a feasible speed for the roughly homogeneous surface. Deterministic MPC completed the maneuver while adhering closely to the desired path (the lane's centerline). The nominal dynamics modeled snow with a coefficient of friction $\mu = 0.25$. Next, Contingency MPC completed the same maneuver while anticipating the potential for ice ($\mu = 0.10$) to cover the road. The red trajectory illustrates the contingency plan that CMPC computed for

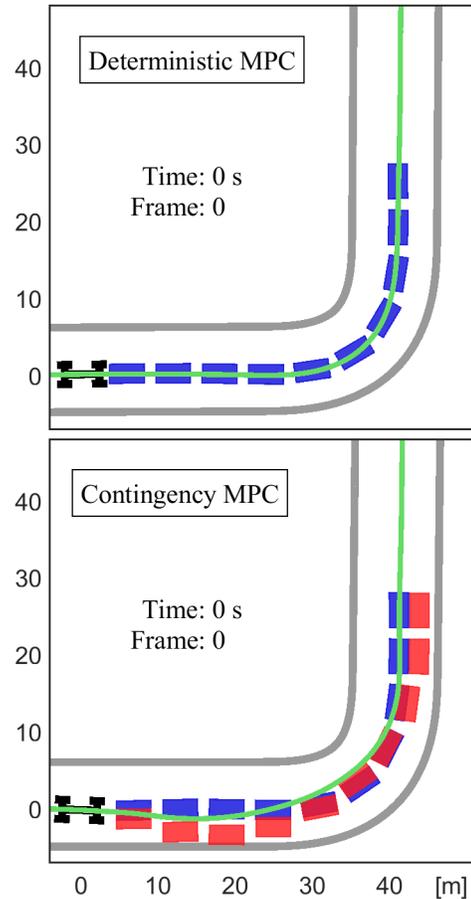


Fig. 5. Deterministic MPC and Contingency MPC experimentally navigated a snow-covered road. Closed-loop behavior in green; nominal horizon in blue; contingency horizon in red. CMPC took a more conservative route. For clarity, Only ten of fifty horizon states are plotted from each trajectory.

this time-step. As the vehicle proceeded under CMPC, it deviated from the desired path to keep the contingency trajectory within path boundaries. It shifted wide, then steered early to make the turn more gentle and conservative. After completing the turn, CMPC returned to the path. In this case the vehicle did not encounter ice and the minor deviation was ultimately unnecessary.

In Fig. 6, the vehicle turned left with the same geometry onto a polished ice surface (cyan). Deterministic MPC's performance is shown in the figure's first pane. Unaware of

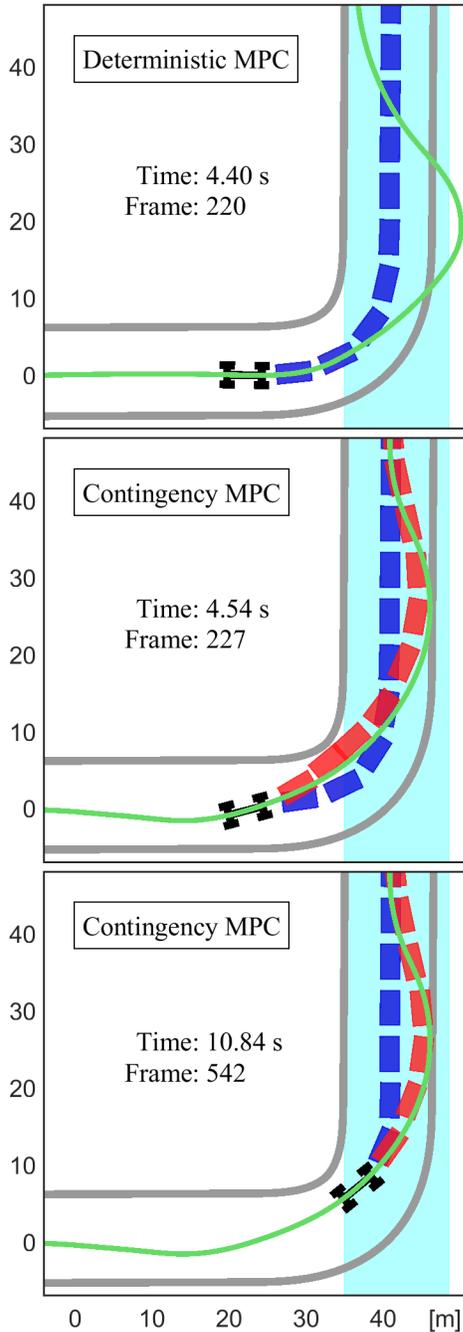


Fig. 6. Deterministic MPC and Contingency MPC experimentally navigated an ice-covered left turn. DMPC did not anticipate ice and lost control. CMPC subtly deviated from the nominal path to increase robustness.

the ice ahead, it naively followed the centerline and could not exert the lateral force necessary to track the path. It ran through the boundary (grey), and recovered after sliding beyond the ice onto snow where it regained traction. After encountering the ice, no avoidance maneuver could have ensured safety. To escape, the vehicle must have behaved differently *before* the emergency.

In panes 2 and 3 of Fig. 6, Contingency MPC demonstrated its effectiveness by repeating the Fig. 5 approach and intuitively deviating from the desired path to design a safer trajectory. The closed-loop behavior resembles the outside-inside-outside cornering technique described by professional racecar driver Taruffi, to maximize speed by minimizing curvature [21]. In this experiment the goal was not to race; however the vehicle's speed profile has been designed to push the scenario's limit. In order to cope with this speed (set a priori), CMPC calculated steering commands to increase the radius of curvature and decrease the tire-force necessary to stay inside the boundary. Thus, the consequence of encountering the extremely slick zamboni-polished surface was significantly mitigated by CMPC, compared to deterministic MPC. In closed-loop (green), CMPC drifted only 20 cm beyond the road edge as it exited the turn.

In Fig. 7, velocity state data is presented from the CMPC test on ice, corresponding to the same time-step as Fig. 6, pane 2. The parallelograms illustrate the stability envelopes from (5d), inside which CMPC attempted to confine the vehicle's state. Measured velocity states are plotted in green, which show the car's actual closed-loop behavior differed significantly from either plan. As it proceeded through the corner, CMPC evaluated the prospect of breaching the stability and environmental envelopes. According to the cost function weights defined in (5b), CMPC allowed a greater

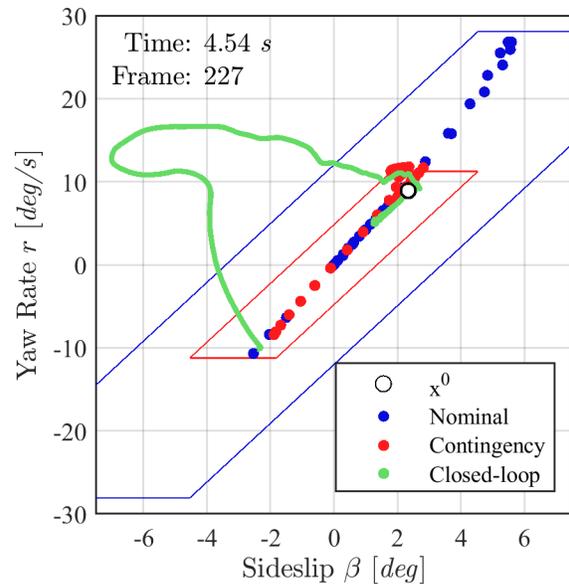


Fig. 7. CMPC prediction horizons and closed-loop trajectory on the $r - \beta$ phase plane. Stability envelopes for snow (blue) and ice (red) were violated in closed-loop to prioritize staying within the road boundary.

violation of the stability guarantee as it prioritized avoidance of the road edge. Setting σ_{env} ten times greater than σ_{stab} influenced CMPC to trade the possibility of losing stability for a minimal collision.

The optimization for this experiment was solved using FORCES Pro on a single i7 processor [22]. The fifty stage deterministic MPC problem converged in an average of 4.5 ms, while Contingency MPC extended average solution times to 11 ms and up to 16 ms in some iterations. This experiment exhibited Contingency MPC's use in a dynamics-based scenario. The prospect of ice on the roadway was encoded into the contingency horizon, which caused controller to act robustly in anticipation of decreased friction.

V. DISCUSSION

The snow and ice experiments successfully demonstrated the concept and efficacy of Contingency MPC. Furthermore, the extreme polished-ice circumstance pushed this convex-linear implementation to its limit, providing insight into how CMPC and similar controllers may be improved in the future. One such limitation was the controller's focus on steering. Indeed, the simplest solution to the presented situation may have been to slow down. Future work could incorporate longitudinal and lateral inputs to CMPC by moving to a nonlinear optimization or by decoupling F_x and F_y .

Furthermore, linearizing the low friction tire model presented difficulty and contributed to the deviation illustrated in Fig. 7. As shown in Fig. 4, the generous linear region present on asphalt nearly vanishes as friction decreases. As a result, linearizations of the low friction model lose accuracy quickly as one moves away from the operating point, allowing some CMPC iterations to predict greater force than was physically available. This problem may be abated by a strategy which maintains greater fidelity of the nonlinear model.

VI. CONCLUSION

In this article, Contingency MPC has been established as a credible strategy to augment a deterministic model predictive controller with robustness. In systems where potential emergency situations can be identified, CMPC maintains an avoidance trajectory while following a desired path. Experimentally, the controller successfully mitigated the effect of an abrupt decrease in road surface friction by consciously sacrificing path tracking performance to increase robustness. It leveraged the knowledge that ice may appear *somewhere* in the environment, and thus planned trajectories which were prepared for its occurrence.

Several avenues appear promising for future development: Applying the framework to additional scenarios may shed valuable insight into the character and range of CMPC. For example, pedestrians may cross unexpectedly or a vehicle ahead may stop abruptly. To investigate its use within a larger system, CMPC could be integrated with an online emergency recognition program to identify contingencies in real-time. In more complex situations, several contingency events may be simultaneously probable. For these cases, CMPC could be augmented to maintain two or more contingency horizons.

ACKNOWLEDGMENT

Thank you Volkswagen Group Research and Engineering Research Lab for experimental support, and Embotech GmbH for academic licensing. Alsterda is supported by Renault Group and the U.S. Dept. of Veterans Affairs G.I. Bill. Brown is supported by Renault Group and the Graduate Research Fellowship from the National Science Foundation.

REFERENCES

- [1] D. Q. Mayne, "Model predictive control: recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, 2014.
- [2] D. Q. Mayne, "Robust and stochastic model predictive control: are we going in the right direction?," *Annual Reviews in Control*, vol. 41, pp. 184–192, 2016.
- [3] P. J. Campo and M. Morari, "Robust model predictive control," in *Proc. American Control Conference*, 1987.
- [4] H. Sartipizadeh and T. L. Vincent, "Uncertainty characterization for robust MPC using an approximate convex hull method," in *Proc. American Control Conference*, 2016.
- [5] D. F. He, H. Huang, and Q. X. Chen, "Quasi-min-max MPC for constrained nonlinear systems with guaranteed input-to-state stability," *Journal of the Franklin Institute*, vol. 351, no. 6, pp. 3405–3423, 2014.
- [6] F. Blanchini, "Control synthesis for discrete time systems with control and state bounds in the presence of disturbances," *Journal of Optimization Theory and Applications*, vol. 65, no. 1, p. 2940, 1990.
- [7] M. Kögel and R. Findeisen, "Robust output feedback MPC for uncertain linear systems with reduced conservatism," in *Proc. 20th IFAC World Congress*, 2017.
- [8] H. Eliasi, M. B. Menhaj, and H. Davilu, "Robust nonlinear model predictive control for a PWR nuclear power plant," *Progress in Nuclear Energy*, vol. 54, no. 1, pp. 177–185, 2012.
- [9] M. B. Saltik, L. Özkan, J. H. Ludlage, S. Weiland, and P. M. Van den Hof, "An outlook on robust model predictive control algorithms: reflections on performance and computational aspects," *Journal of Process Control*, vol. 61, pp. 77–102, 2018.
- [10] M. Farina, L. Giulioni, and R. Scattolini, "Stochastic linear model predictive control with chance constraints - a review," *Journal of Process Control*, vol. 44, pp. 53–67, 2016.
- [11] S. Lucia, T. Finkler, and S. Engell, "Multi-stage nonlinear model predictive control applied to a semi-batch polymerization reactor under uncertainty," *Journal of Process Control*, 2013.
- [12] D. Krishnamoorthy, E. Suwartadi, B. Foss, S. Skogestad, and J. Jaschke, "Improving scenario decomposition for multistage MPC using a sensitivity-based path-following algorithm," *IEEE Control Systems Letters*, vol. 2, no. 4, pp. 581 – 586, 2018.
- [13] J. Richalet, A. Rault, J. L. Testud, and J. Papon, "Model predictive heuristic control: applications to industrial processes," *Automatica*, vol. 14, no. 5, pp. 413–428, 1978.
- [14] M. Brown, J. Funke, S. Ertlien, and J. C. Gerdes, "Safe driving envelopes for path tracking in autonomous vehicles," *Control Engineering Practice*, vol. 61, pp. 307–316, 2017.
- [15] V. Zhang, S. M. Thornton, and J. C. Gerdes, "Tire modeling to enable model predictive control of automated vehicles from standstill to the limits of handling," in *Proc. 14th International Symposium on Advanced Vehicle Control*, 2018.
- [16] S. M. Ertlien, J. Funke, and J. C. Gerdes, "Incorporating non-linear tire dynamics into a convex approach to shared steering control," in *Proc. American Control Conference*, 2014.
- [17] H. Pacejka, *Tire and Vehicle Dynamics*. Butterworth-Heinemann, 3rd ed., 2012.
- [18] J. Funke, M. Brown, S. M. Ertlien, and J. C. Gerdes, "Collision avoidance and stabilization for autonomous vehicles in emergency scenarios," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 4, pp. 1204 – 1216, 2017.
- [19] C. G. Bobier, S. Joe, and J. C. Gerdes, "Sliding surface envelope control: keeping the vehicle within a safe state-space boundary," in *Proc. Dynamic Systems and Control Conference*, 2011.
- [20] C. E. Beal and J. C. Gerdes, "Model predictive control for vehicle stabilization at the limits of handling," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 4, pp. 1258 – 1269, 2013.
- [21] P. Taruffi, *The Technique of Motor Racing*. Bentley Publishers, 2003.
- [22] A. Domahidi and J. Jerez, "FORCES Professional," 2014.