MITSUBISHI ELECTRIC RESEARCH LABORATORIES
https://www.merl.com

Cornering Stiffness Adaptive, Stochastic Nonlinear Model Predictive Control for Vehicles

Vaskov, Sean; Berntorp, Karl; Quirynen, Rien;

TR2021-056 June 04, 2021

Abstract

The vehicle control behavior is highly dependenton the road surface. However, accurate and precise models for the tire–road interaction are typically unknown a priori. It is therefore important that the vehicle's control algorithm updates its tire-force model, to adapt to the changing conditions. In this paper, we propose a stochastic nonlinear model-predictive control (SNMPC) scheme that uses a linear tire-force model, where the mean and covariance of the cornering stiffness parameters are estimated and updated online. We formulate constraints based on the stiffness estimates to ensure that the vehicle maintains stability on low-friction surfaces. In extensive simulations, where the road surface transitions from asphalt to snow, we compare the proposed controller with various MPC implementations; for example, the proposed approach reduces average closed-loop cost over 30% on aggressive maneuvers, when compared to a non-stochastic controller.

American Control Conference (ACC) 2021

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

Cornering Stiffness Adaptive, Stochastic Nonlinear Model Predictive Control for Vehicles

Sean Vaskov[1,2], Karl Berntorp[1], Rien Quirynen [1]

***Abstract*— The vehicle control behavior is highly dependent on the road surface. However, accurate and precise models for the tire–road interaction are typically unknown a priori. It is therefore important that the vehicle's control algorithm updates its tire-force model, to adapt to the changing conditions. In this paper, we propose a stochastic nonlinear model-predictive control (SNMPC) scheme that uses a linear tire-force model, where the mean and covariance of the cornering stiffness parameters are estimated and updated online. We formulate constraints based on the stiffness estimates to ensure that the vehicle maintains stability on low-friction surfaces. In extensive simulations, where the road surface transitions from asphalt to snow, we compare the proposed controller with various MPC implementations; for example, the proposed approach reduces average closed-loop cost over 30% on aggressive maneuvers, when compared to a non-stochastic controller.**

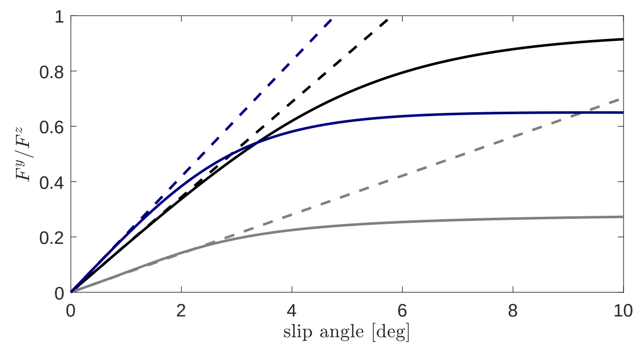

Fig. 1. Examples of normalized lateral force as a function of slip angle α for dry asphalt (black), wet asphalt (navy), and snow (gray). The solid lines show the tire-force curves defined by a Pacejka model. The dashed lines show a linear tire model where the slope is calculated at 0 degrees slip.

I. INTRODUCTION

Control systems for autonomous vehicles actuate the vehicle through tire–road contact; therefore knowledge of the tire–road relation is of high importance. The interaction between tire and road is highly nonlinear, and the parameters describing the nonlinear relation vary heavily based on the road surface and other tire properties [1], [2]. Figure 1 shows examples of the tire-force variation with the wheel slip for three different surfaces. The force-slip relation is approximately linear for small slip values, which are typical when driving in normal conditions. Knowledge of the tire stiffness can be used directly in ADAS [3], [4], and even partial knowledge of the tire stiffness can be used to classify surface types for road-condition monitoring [2], [5].

Model Predictive Control (MPC) has been effective in several automotive applications [6]–[8]. MPC solves an optimization problem, where a dynamic model of the vehicle is integrated over a fixed time horizon to minimize a user-specified cost subject to constraints on the inputs and states. In many ADAS applications, this often leads to a nonlinear MPC (NMPC) problem due to the vehicle model and constraints. For an overview of integration schemes with sensitivity analysis to treat explicit and implicit nonlinear differential equations in embedded NMPC, see [9]. In sequential quadratic programming (SQP) based NMPC, a tailored convex solver is used to solve a sequence of structured quadratic programs (QPs) [10]. In recent years, many such algorithms have been developed to exploit particular sparsity structures that arise in SQP based NMPC, such as the recently proposed QP solver in [11] and references therein. In [12], an optimization algorithm is proposed for stochastic NMPC (SNMPC), which uses a tailored Jacobian approximation along with an adjoint-based SQP method.

Since the performance of MPC depends heavily on an accurate model, recent studies have focused on adaptive controllers, where uncertain parameters are estimated and the model is updated online. In [13], a robust MPC formulation is proposed, where a parameter associated with the steering offset is estimated. In [14], [15], least-squares algorithms are used to estimate the cornering stiffness and road friction, which are utilized in the MPC model and constraints. These two works do not consider uncertainty of the estimated stiffness, and [13]–[15] all use linear vehicle models.

In [16], a particle-filter based algorithm is proposed, which estimates the mean and covariance of the tire stiffness using data from commonly available inertial sensors. Prior work, see [3], utilized this cornering-stiffness estimator in NMPC by selecting from a library of predefined nonlinear tire models. However, relying on a fixed model for surfaces with large variability (e.g., packed vs loose snow) may result in poor controller performance. Furthermore, it is not obvious how to incorporate uncertainty associated with the stiffness estimate into a library-based approach. In the present paper, we propose an SNMPC that uses a linear tire-force model. We directly use the mean and covariance from the stiffness estimator to approximate chance constraints in the optimal control problem. The contributions of this paper are that we:

1) Incorporate the stiffness estimate and uncertainty from the stiffness estimator presented in [16] into the SNMPC problem formulation proposed in [12].
2) Develop a set of stability constraints, dependent on the stiffness estimate, that enable the controller to perform moderately aggressive maneuvers.

[1]Mitsubishi Electric Research Laboratories, Cambridge, MA, 02139, USA. Email: karl.o.berntorp@ieee.org
[2]University of Michigan, Ann Arbor, MI, 48109, USA.

3) Demonstrate through simulations that the proposed formulation improves robustness and performance over non-stochastic and non-adaptive controllers.

II. VEHICLE MODELING

We use a single-track chassis model that includes the longitudinal velocity v^x, lateral velocity v^y, yaw rate $\dot{\psi}$, and wheel angle δ as states. The inputs to the vehicle model are the front and rear wheel speeds ω_f, ω_r and the tire-wheel angle rate of change $\dot{\delta}$. As shown in [17], a single-track model is sufficiently accurate where the tire forces reach the nonlinear region but the maneuvers are not aggressive enough to result in large roll angles. The single-track model lumps together the left and right wheel on each axle, and roll and pitch dynamics are neglected. Thus, the model has two translational and one rotational degrees of freedom. The model dynamics read as

$$\begin{bmatrix} \dot{v}^x \\ \dot{v}^y \\ \ddot{\psi} \end{bmatrix} = \begin{bmatrix} \frac{1}{m}(F_f^x \cos(\delta) + F_r^x - F_f^y \sin(\delta)) + v_y \dot{\psi} \\ \frac{1}{m}(F_f^y \cos(\delta) + F_r^y + F_f^x \sin(\delta)) - v_x \dot{\psi} \\ \frac{1}{I}(l_f F_f^y \cos(\delta) - l_r F_r^y + l_f F_f^x \sin(\delta)) \end{bmatrix}, \tag{1}$$

where F^x, F^y are the longitudinal/lateral tire forces and the subscripts f, r stand for front and rear, respectively, m is the vehicle mass, I is the vehicle inertia about the vertical axis, δ is the front-wheel steering angle, and l_f and l_r are the distance from the front and rear axles to the center of mass. The normal force F_i^z resting on each front/rear wheel are approximated as

$$F_f^z = mg\frac{l_r}{l}, \quad F_r^z = mg\frac{l_f}{l}, \tag{2}$$

where the wheel base is $l = l_f + l_r$. The slip angles α_i and slip ratios λ_i are defined as in [18], [19],

$$\alpha_i = -\arctan\left(\frac{v_i^y}{v_i^x}\right), \quad \lambda_i \quad = \frac{R_w\omega_i - v_i^x}{\max(R_w\omega_i, v_i^x)}, \tag{3}$$

where $i \in \{f, r\}$ and R_w is the wheel radius, and v_i^x and v_i^y are the longitudinal and lateral wheel velocities for wheel i with respect to an inertial system, expressed in the coordinate system of the wheel. The tire forces are computed with the Magic Formula model [18], and combined loading is based on the friction ellipse as follows

$$\begin{aligned} &F_i^x = \mu_i^x F_i^z \sin(D_i^x \arctan(B_i^x(1 - E_i^x)\lambda_i + E_i^x \arctan(B_i^x \lambda_i))), \\ &F_i^y = \eta_i \mu_i^y F_i^z \sin(D_i^y \arctan(B_i^y(1 - E_i^y)\alpha_i + E_i^y \arctan(B_i^y \alpha_i))), \\ &\eta_i = \sqrt{1 - \left(\frac{F_i^x}{\mu_i^x F_i^z}\right)^2}, \end{aligned} \tag{4}$$

where μ_i^j, B_i^j, D_i^j and E_i^j, for $i \in \{f, r\}, j \in \{x, y\}$, are the friction coefficients and stiffness, shape, and curvature factors. In (4), the longitudinal force does not explicitly depend on the lateral slip, and it is possible to use more accurate models to represent the combined slip [17], [18]. Pacejka's magic formula (4) exhibits the typical saturation behavior in the tire forces as illustrated in Figure 1.

III. CORNERING STIFFNESS ESTIMATION

The tire-stiffness estimator is based on a recently developed adaptive particle-filter approach, see [16]. An important feature of the estimator is that it only relies on sensors commonly available in production vehicles. The method employs the single-track vehicle model (1) and a linear approximation of the front and rear tire forces,

$$F_i^x \approx C_i^x \lambda_i, \quad F_i^y \approx C_i^y \alpha_i, \tag{5}$$

where C_i^x and C_i^y are the longitudinal and lateral stiffness, respectively. As seen in Figure 1, the linear approximation is valid at low slip values. In this work, we only estimate the lateral cornering stiffness and assume $C_i^x \approx 2C_i^y$. This approximation provides a coarse update of the longitudinal stiffness without introducing additional parameters into the tire-force estimator. The linear relationship was chosen based on the models used in simulation and could alternatively be fit with experimental data. Since the maneuvers in this work will not require large longitudinal accelerations, accurately modelling the longitudinal stiffness is not critical to the controller performance. However, the estimation algorithm could estimate both the lateral and longitudinal forces, if desired.

The stiffness values in (5) are decomposed into a nominal and unknown part,

$$C_i^y = C_{i,n}^y + w_{i,k}, \tag{6}$$

where $C_{i,n}^y$ is the nominal value of the cornering stiffness, for example, a priori determined on a nominal surface, and $w_{i,k}$ is a time-varying, unknown part. We model the unknown stiffness components as random process noise $\boldsymbol{w}_k \in \mathbb{R}^{n_w}$ acting on the otherwise deterministic system. The noise is assumed Gaussian distributed according to $\boldsymbol{w}_k \sim \mathcal{N}(\Delta \boldsymbol{C}_k, \boldsymbol{\Sigma}_k)$, where $\Delta \boldsymbol{C}_k$ and $\boldsymbol{\Sigma}_k$ are the unknown, usually time varying, mean and covariance. Inserting (5)–(6) into (1) and discretizing using forward-Euler with a sampling period t_s gives the discrete-time dynamics of the form

$$\boldsymbol{x}_{k+1} = \boldsymbol{x}_k + t_s \boldsymbol{f}(\boldsymbol{x}_k, \boldsymbol{u}_k) + t_s \boldsymbol{g}(\boldsymbol{x}_k, \boldsymbol{u}_k)\boldsymbol{w}_k, \tag{7}$$

$$\boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}) = \begin{bmatrix} \frac{C_{f,n}^y}{m}(2\lambda_f \sin\delta + \alpha_f \cos\delta) + \frac{C_{r,n}^y}{m}\alpha_r - v^x \dot{\psi} \\ \frac{C_{f,n}^y l_f}{I}(2\lambda_f \sin\delta + \alpha_f \cos\delta) - \frac{C_{r,n}^y l_r}{I}\alpha_r \end{bmatrix}, \tag{8}$$

$$\boldsymbol{g}(\boldsymbol{x}, \boldsymbol{u}) = \begin{bmatrix} \frac{1}{m}(2\lambda_f \sin\delta + \alpha_f \cos\delta) & \frac{1}{m}\alpha_r \\ \frac{l_f}{I}(2\lambda_f \sin\delta + \alpha_f \cos\delta) & \frac{-l_r}{I}\alpha_r \end{bmatrix}, \tag{9}$$

where the subscript k refers to the current timestep. The estimator uses the lateral (and optionally longitudinal) acceleration and yaw-rate measurements and models the bias $\boldsymbol{b}_k$ of the inertial measurements as a random walk, which results in a measurement model

$$\boldsymbol{y}_k = \boldsymbol{h}(\boldsymbol{x}_k, \boldsymbol{u}_k) + \boldsymbol{b}_k + \boldsymbol{d}(\boldsymbol{x}_k, \boldsymbol{u}_k)\boldsymbol{w}_k + \boldsymbol{e}_k, \tag{10}$$

$$\boldsymbol{h}(\boldsymbol{x}, \boldsymbol{u}) = \begin{bmatrix} \frac{C_{f,n}^y}{m}(2\lambda_f \sin\delta + \alpha_f \cos\delta) + \frac{C_{r,n}^y}{m}\alpha_r \\ \dot{\psi} \end{bmatrix}, \tag{11}$$

$$\boldsymbol{d}(\boldsymbol{x}, \boldsymbol{u}) = \begin{bmatrix} \frac{1}{m}(2\lambda_f \sin\delta + \alpha_f \cos\delta) & \frac{1}{m}\alpha_r \\ 0 & 0 \end{bmatrix}. \tag{12}$$

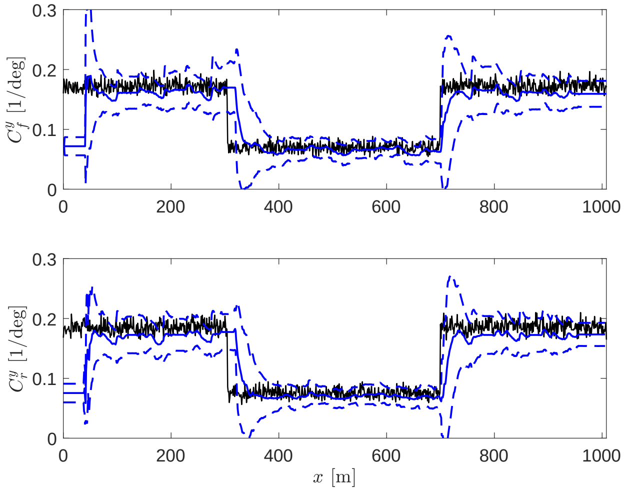

Fig. 2. Stiffness estimates for a surface switching from dry asphalt to snow and back. The black line is the true stiffness, the slope of the nonlinear tire-force curve at $\alpha_i = 0$. The blue solid and dashed lines are the mean and 95% confidence interval from the stiffness estimator. The true stiffness is underestimated when the tires saturate.

Each particle of the estimator contains, in addition to the state vector, an estimated mean value $\Delta \boldsymbol{C}_k$ of the stiffness noise and the corresponding covariance estimate $\boldsymbol{\Sigma}_k$. In this work, we use the weighted average of the mean and covariance among all particles, but different types of risk metrics could be considered in future work. Note that because of the inertial sensor measurements, the stiffness components enter both in the vehicle model and the measurement model through $\boldsymbol{w}_k$, which implies that the estimation model has a dependence between the process and measurement noise.

Remark 1: Due to the approximation in (5), the stiffness estimator operates under the assumption of moderate steering angles and driving/braking torques. In the implementation, the estimator is activated only when the wheel angle and slip ratios are within a predefined threshold. Additionally, the estimator is deactivated when the wheel angle is near zero since the system becomes unobservable [16].

Figure 2 shows the output from the stiffness estimator on a surface switching from dry asphalt to snow and back. The estimator uses a sampling period of $t_s = 0.01$ s.

The "true stiffness" is defined as the slope of the tire-force curve at $\alpha_i = 0$. In Figure 2, we can see that the true stiffness is underestimated at times on both surfaces, as a result of tire saturation. Figure 3 provides a simple illustration of why this occurs for an asphalt tire model. When the vehicle is operating at nonzero slip angles, the estimated tire-force model can be thought of as a line between the origin and the true tire force. The slope of this line (the estimated stiffness) decreases as the slip angle increases and the tire-force curve flattens.

IV. STOCHASTIC NMPC FORMULATION

Let us consider nonlinear systems of the form

$$\boldsymbol{x}_{k+1} = \boldsymbol{f}(\boldsymbol{x}_k, \boldsymbol{u}_k, \boldsymbol{w}_k), \tag{13}$$

where $\boldsymbol{x}_k \in \mathbb{R}^{n_x}$ denotes the state, $\boldsymbol{u}_k \in \mathbb{R}^{n_u}$ the control inputs, $\boldsymbol{w}_k \in \mathbb{R}^{n_w}$ the process noise, and $\boldsymbol{f} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_w} \to \mathbb{R}^{n_x}$ the system dynamics. In this context, the system dynamics are given by (1), appended with differential equations for the position, heading, and front-wheel angle. The tire-force equations are linear, as in (5), with the stiffness updated online from the estimator described in Section III. The disturbance $\boldsymbol{w}_k \sim \mathcal{N}(\Delta C, \boldsymbol{\Sigma})$ is assumed to be normally distributed with mean ΔC, covariance matrix $\boldsymbol{\Sigma}$ updated from the stiffness estimator.

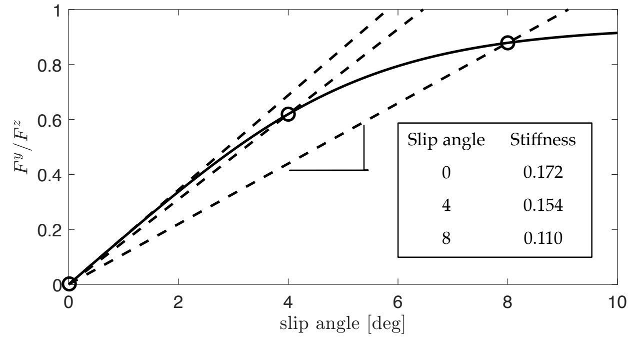

Fig. 3. Linear approximation of the tire forces at slip angles of 0, 4, 8 deg for an asphalt tire model. The solid line is the true tire-force curve, while dashed lines are the linear approximations. The stiffness (slope of the line) decreases as the slip angle increases.

At each sampling time, based on the current state estimate $\hat{\boldsymbol{x}}_t$ and covariance $\boldsymbol{P}_t$, the SNMPC solves

$$\begin{aligned}
\min_{\mathbf{x},\mathbf{u},\boldsymbol{P}} \quad & \sum_{k=0}^{N-1} l(\boldsymbol{x}_k, \boldsymbol{u}_k) \\
\text{s.t.} \quad & \begin{cases} \forall k \in \{0, \dots, N-1\}, \\ 0 = \boldsymbol{x}_{k+1} - \boldsymbol{f}(\boldsymbol{x}_k, \tilde{\boldsymbol{u}}_k, \Delta C_t), \\ 0 = \boldsymbol{x}_0 - \hat{\boldsymbol{x}}_t, \\ \boldsymbol{P}_{k+1} = \boldsymbol{A}_k \boldsymbol{P}_k \boldsymbol{A}_k^\top + \boldsymbol{B}_k \boldsymbol{\Sigma}_t \boldsymbol{B}_k^\top, \ \boldsymbol{P}_0 = \boldsymbol{P}_t, \\ Pr\left(\mathbf{c}(\boldsymbol{x}_k, \boldsymbol{u}_k) \leq 0\right) \geq 1 - \epsilon, \end{cases}
\end{aligned} \tag{14}$$

where the overall control action is in the feedforward-feedback form $\tilde{\boldsymbol{u}}_k = \boldsymbol{u}_{\text{cte}} + K\boldsymbol{x}_k + \boldsymbol{u}_k$ due to a prestabilizing controller, and the Jacobian matrices read as $\boldsymbol{A}_k = \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}}(\boldsymbol{x}_k, \tilde{\boldsymbol{u}}_k, \Delta C_t)$ and $\boldsymbol{B}_k = \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{w}}(\boldsymbol{x}_k, \tilde{\boldsymbol{u}}_k, \Delta C_t)$. The state covariance propagation equations correspond to the extended Kalman filtering (EKF) approach, similar to [20].

A. Objective Function and Inequality Constraints

We consider the stage cost in (14) to be

$$l(\cdot) = \frac{1}{2}\|\boldsymbol{x}_k(\cdot) - \boldsymbol{x}_{\text{ref},k}\|_Q^2 + \frac{1}{2}\|\tilde{\boldsymbol{u}}_k(\cdot) - \tilde{\boldsymbol{u}}_{\text{ref},k}\|_R^2. \tag{15}$$

We enforce the following constraints $\mathbf{c}(\boldsymbol{x}_k, \boldsymbol{u}_k) \leq 0$ in the optimal-control problem of (14):

$$y_{\text{min,k}} \leq y_k \leq y_{\text{max,k}}, \tag{16a}$$

$$|\delta_k| \leq \delta_{\text{max}}, \quad |\dot{\delta}_k| \leq \dot{\delta}_{\text{max}}, \tag{16b}$$

$$|\lambda_{i,k}| \leq \lambda_{\text{max}}, \ i \in \{f, r\}, \tag{16c}$$

$$|\dot{\psi}_k v_k^x| \leq 0.85\mu g, \quad \left|\frac{v_k^y}{v_k^x}\right| \leq \tan^{-1}(0.02\mu g). \tag{16d}$$

Eq. (16a) bounds the lateral position, and is used to ensure that the vehicle stays on the road. Obstacle avoidance constraints could be considered in future work. Eqs. (16b)-(16c) bound the wheel angle, wheel angle rate, and slip ratios. The constraints in (16d) prevent the vehicle from entering regions of high lateral acceleration and side slip, and can be found in [19, Chapter 8]. We refer to (16d) as *stability constraints*.

The stability constraints depend on the road friction μ, a parameter whose estimation is widely studied [21]. Experimental studies suggest that using a monotonic relationship is sufficient to differentiate between asphalt and snow [2], [5]. In this work, we use a linear relationship to approximate the road friction as a function of the cornering stiffness estimate,

$$\mu \approx \min\left(\frac{a(C^y_{f,n} + \Delta C^y_f + C^y_{r,n} + \Delta C^y_r)}{2}, 1\right), \qquad (17)$$

where a is a constant that was fit from Pacejka models for asphalt and snow. This relationship proved to be effective in our simulations; finding an optimal relationship to use could be the subject of future work. The central idea of (17) is that the bounds on the acceleration and sideslip should tighten as the road friction, and consequently the cornering stiffness, decreases. For surfaces such as wet asphalt, which may have a high cornering stiffness but lower road friction, (17) is conservative because the stiffness estimator underestimates the true stiffness as the tires saturate (as in Figure 3).

B. Probabilistic Chance Constraints

To enforce the probabilistic chance constraints in (14), we reformulate them as deterministic constraints as in [20], where the j^{th} constraint is written as

$$c_j(\boldsymbol{x}_k, \boldsymbol{u}_k) + \nu\sqrt{\frac{\partial c_j}{\partial \boldsymbol{x}_k}\boldsymbol{P}_k\frac{\partial c_j}{\partial \boldsymbol{x}_k}^T} \le 0, \qquad (18)$$

where ν is referred to as the back-off coefficient and depends on the desired probability threshold ϵ and assumptions about the resulting state distribution. The backoff coefficient for Cantelli's inequality, $\nu = \sqrt{\frac{1-\epsilon}{\epsilon}}$, holds regardless of the underlying distribution but is conservative. We assume normally-distributed state trajectories and set $\nu = \sqrt{2}\text{erf}^{-1}(1-2\epsilon)$, where $\text{erf}^{-1}(\cdot)$ is the inverse error function.

C. Software Implementation Aspects

We use the SNMPC implementation that was proposed recently in [12], based on an SQP optimization algorithm in which a series of QP approximations are solved using the PRESAS QP solver [11]. The algorithm uses a tailored Jacobian approximation along with an adjoint-based SQP method that allows for the numerical elimination of the covariance matrices from the SQP subproblem, which reduces the computation time when compared to standard SQP formulations for SNMPC [12]. Note that one SQP iteration per control time step is typically performed for real-time implementations of NMPC, as discussed in [10].

D. Illustrative Example for SNMPC Formulation

We set up a control reference that intentionally violates the lower constraint on the lateral position, to illustrate how the chance constraint is respected when we solve the SNMPC optimization problem in (14). The mean cornering stiffness values correspond to a snow surface. We set the standard deviation to be roughly 10% of the mean values. The least-squares cost (15) prioritizes the lateral position and wheel speed inputs. A timestep of $t_s = 0.05$ s with a prediction horizon of 2 s is used. The solution trajectory is shown in Figure 4. We integrate the dynamic model forward for 1e5 disturbance realizations, and see that the chance constraint for $\epsilon = 0.05$ is approximated to within 1%.

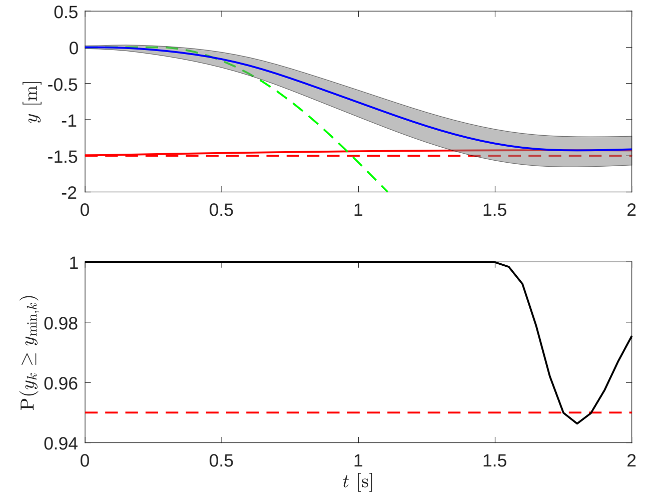

Fig. 4. Illustration of a chance constraint enforced on the lateral position. In the top plot, the green dashed line is the reference trajectory. The red dashed line is the lateral constraint, and the red solid line is the tightened, chance constraint. The blue line is the mean trajectory from solving (14). The gray shaded region indicates the upper and lower bounds for 1e5 random disturbance realizations. The lower plot shows the probability of constraint satisfaction of the 1e5 simulations (black) and the desired threshold (red). The chance constraint for $\epsilon = 0.05$ is approximated to within 1%

V. SIMULATION RESULTS

The first case study requires the vehicle to track nine double lane-change and return maneuvers, with the middle three on snow and the rest on dry asphalt. To investigate the learning behavior of the controller, the surface change occurs during a straight portion, where the stiffness is unobservable. The spacing of the asphalt maneuvers is similar to IS0 3888-2 [22]; whereas the snow maneuvers are elongated for feasibility. The reference velocity is fixed to 17 m/s. The reference is generated with Bezier polynomials and the position, heading, longitudinal velocity, and yaw rate are given to the controllers to track. The lateral constraints we enforce are that the vehicle is not allowed to leave the road boundaries. The simulation model uses the Pacejka tire model described in Section II. The Pacejka parameters for each road surface are randomly perturbed at each controller timestep, with samples drawn from a uniform distribution up to ± 5% for asphalt and 10% for snow.

We compare the following 5 NMPC controllers:

1) STOCHASTIC: proposed SNMPC controller with online adaptation to stiffness-estimation results.
2) ADAPTIVE: nominal NMPC controller with online adaptation to the mean cornering stiffness.
3) SNOW: nominal NMPC with cornering stiffness fixed to snow parameter values.

4) ASPHALT: nominal NMPC with cornering stiffness fixed to dry asphalt parameter values.
5) ORACLE: NMPC with true nonlinear tire-force model.

We include the ORACLE controller to provide a lower bound on cost and constraint violations for the simulations. Its performance cannot be achieved in practice because it is given the exact tire force curve used by the simulation model; in reality there will be model mismatch due to inaccuracies in both the tire force and single-track vehicle models.

All controllers perform 1 SQP iteration per time step [10] and the nominal NMPC controllers 2-5 do not have stochastic constraints. For the stability constraints in (16d), the ASPHALT and SNOW controllers assume road friction values of $\mu = 1.0$ and 0.35, respectively. Since the ORACLE utilizes a nonlinear tire model, the stability constraints (16d) are not enforced. The least-squares cost (15) prioritizes the lateral position and wheel speed inputs. A timestep of $t_s = 0.05$ with a prediction horizon of 2 s is used. The stiffness estimator is run at 100 Hz. The constraint satisfaction probability for the STOCHASTIC controller is set to 95%, i.e., $\epsilon = 0.05$. The metrics we use to evaluate the controllers are cost and score, and are computed as follows:

$$\text{Cost} = \sum_k l(\boldsymbol{x}_k, \boldsymbol{u}_k), \tag{19}$$

$$\text{Score} = \sum_k ((y_k - y_{\max})_+ + (y_{\min} - y_k)_+)\, t_s, \tag{20}$$

where $(\cdot)_+ = \max(\cdot, 0)$. The results of 200 trials are shown in Table I. In most trials, the ASPHALT controller destabilizes the vehicle and the trials were terminated early; the reported cost and score is summed up to the point of termination.

Figure 5 shows the trajectories, and Figure 6 shows the stability constraints in (16d) for an example trial. The ASPHALT controller is unable to safely navigate the maneuvers on snow; whereas the SNOW controller behaves conservatively on asphalt. The STOCHASTIC and ADAPTIVE controllers overshoot the first maneuver on snow, but are able to match the performance of the SNOW and ORACLE controllers once they have learned about the surface change. The average cost for the STOCHASTIC controller is 1% less than the ADAPTIVE controller, 86% less than the SNOW controller, and only 29% more than the ORACLE. The STOCHASTIC controller does not violate the lateral constraints in this example, and performs better on the score metric than the ADAPTIVE controller.

TABLE I
RESULTS FOR 200 RANDOM TRIALS ON DRY ASPHALT/SNOW AT 17 M/S

NMPC Controller	Cost		Score	
	mean	max	mean	max
STOCHASTIC	0.339	0.448	0	0
ADAPTIVE	0.342	0.480	1.4e-4	0.013
SNOW	2.463	2.500	0	0
ASPHALT	136.0	562.8	3.254	20.70
ORACLE	0.263	0.268	0	0

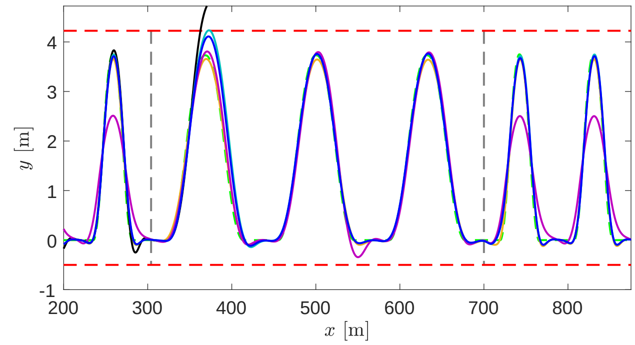

Fig. 5. Position trajectories for a sample trial at 17 m/s where the middle 3 maneuvers are on snow and the others on dry asphalt. Red and green dashed lines are the constraints and reference. The gray dashed lines indicate the surface changes. The blue, cyan, magenta, black, and gold lines indicate the trajectories for the STOCHASTIC, ADAPTIVE, SNOW, ASPHALT, and ORACLE controllers. The STOCHASTIC controller is able to satisfy the lateral constraints and closely match the performance of the ORACLE controller after it learns about the surface change.

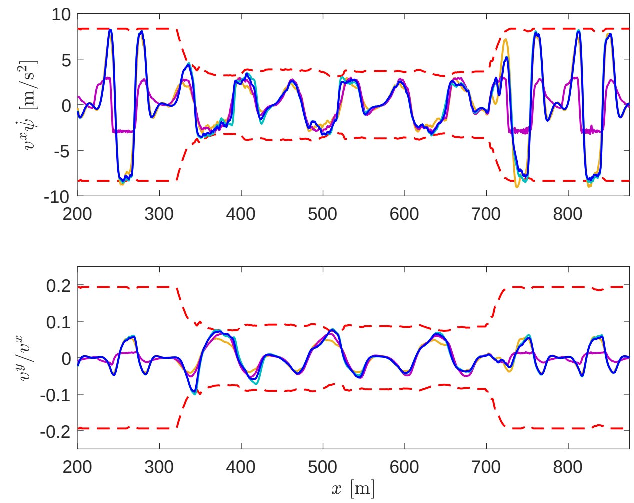

Fig. 6. Stability constraints for the sample trial in Figure 5, where the middle portion is on snow. Red dashed lines are the constraint boundaries, where the road friction is calculated with (17) using the estimator output from the STOCHASTIC controller. The constraints tighten during the snow portion. Coloring for the controllers is the same as in Figure 5. The ASPHALT controller destabilizes the vehicle and is omitted for clarity.

The second case study uses the same setup, except we increase the vehicle speed to 19 m/s. The results of 100 trials are shown in Table II.

Figure 7 shows the trajectories for a sample trial. Compared to the previous case study, all of the controllers have an increased cost and, aside from the ORACLE, some constraint violations. The SNOW controller frequently violates the lateral constraints due to the fact that it is using a linear tire model with fixed stiffness parameters and the tires saturate at the faster velocity. The ADAPTIVE controller violates the lateral constraints frequently during the first snow maneuver, since it does not take uncertainty in the stiffness estimate into account while it is learning the surface change. The average score for the STOCHASTIC controller is 94% less than the ADAPTIVE controller and 96% less than the SNOW controller. The average cost for the STOCHASTIC controller is 34% less than the ADAPTIVE controller, 64% less than the SNOW controller, but now 68% more than the ORACLE. The maximum cost for the ADAPTIVE controller also increases

TABLE II

RESULTS FOR 200 RANDOM TRIALS ON DRY ASPHALT/SNOW AT 19 M/S

NMPC Controller	Cost		Score	
	mean	max	mean	max
STOCHASTIC	1.193	1.733	1.2e-3	0.037
ADAPTIVE	1.814	3.881	0.021	0.086
SNOW	3.329	3.442	0.034	0.102
ASPHALT	219.7	834.8	5.96	16.15
ORACLE	0.710	0.725	0	0

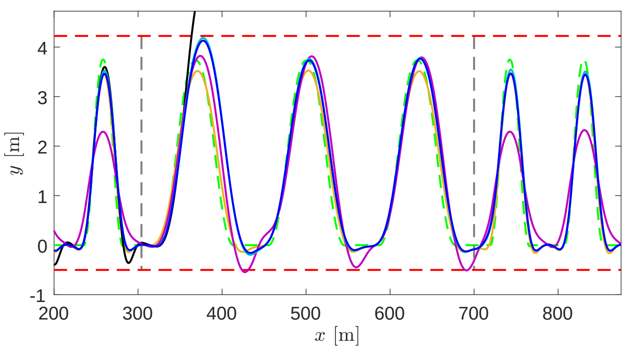

Fig. 7. Position trajectories for a sample trial at 19 m/s where the middle 3 maneuvers are on snow and the others on dry asphalt. Red and green dashed lines are the constraints and reference. The gray dashed lines indicate the surface change. Coloring for the controllers is the same as in Figure 5. The STOCHASTIC controller satisfies the lateral constraints and closely match the performance of the ORACLE after it learns the surface change.

significantly, relative to the STOCHASTIC.

Overall, the results show that the STOCHASTIC controller is able to closely match the performance of the ORACLE controller once it has learned about the road surface. The constraint violations and cost for the STOCHASTIC controller are incurred mainly during the first snow maneuver, as the surface change occurs during a straight portion where the cornering stiffness is not observable. Improving the design of the reference trajectory to encourage persistent excitation, modulating the forgetting factor in the estimator, or incorporating a road friction forecast based on external sensors could greatly improve the performance of the STOCHASTIC and ADAPTIVE controllers.

VI. CONCLUSION

This paper presents an adaptive SNMPC formulation for vehicle control, that uses a linear tire-force model, where the mean and covariance of the cornering stiffnesses are estimated online with a particle-filter based approach. We enforce chance constraints on the inputs, lateral position, lateral acceleration and side slip; the bounds for the latter are varied based on the stiffness estimate. Simulation results show that on moderately aggressive maneuvers, with surfaces varying between dry asphalt and snow, the proposed formulation outperforms controllers with fixed stiffness parameters. Additionally, the proposed approach has fewer constraint violations and lower cost when compared to an adaptive controller that does not incorporate the estimation uncertainty, and it achieves comparable performance to an oracle controller that is given the true nonlinear tire models of the simulated surface.

REFERENCES

[1] J. Svendenius, "Tire modeling and friction estimation," Ph.D. dissertation, Department of Automatic Control, Lund University Lund, Sweden, 2007.

[2] F. Gustafsson, "Slip-based tire-road friction estimation," *Automatica*, vol. 33, no. 6, pp. 1087–1099, 1997.

[3] K. Berntorp, R. Quirynen, T. Uno, and S. Di Cairano, "Trajectory tracking for autonomous vehicles on varying road surfaces by friction-adaptive nonlinear model predictive control," *Vehicle System Dynamics*, vol. 58, no. 5, pp. 705–725, 2019.

[4] S. Di Cairano, H. E. Tseng, D. Bernardini, and A. Bemporad, "Vehicle yaw stability control by coordinated active front steering and differential braking in the tire sideslip angles domain," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 4, pp. 1236–1248, 2012.

[5] C. Ahn, H. Peng, and H. E. Tseng, "Robust estimation of road friction coefficient using lateral and longitudinal vehicle dynamics," *Vehicle System Dynamics*, vol. 50, no. 6, pp. 961–985, 2012.

[6] P. Falcone, F. Borrelli, J. Asgari, H. E. Tseng, and D. Hrovat, "Predictive active steering control for autonomous vehicle systems," *IEEE Transactions on control systems technology*, vol. 15, no. 3, pp. 566–580, 2007.

[7] S. Di Cairano, U. Kalabić, and K. Berntorp, "Vehicle tracking control on piecewise-clothoidal trajectories by mpc with guaranteed error bounds," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 709–714.

[8] R. Quirynen, K. Berntorp, and S. Di Cairano, "Embedded optimization algorithms for steering in autonomous vehicles based on nonlinear model predictive control," in *2018 Annual American Control Conference (ACC)*. IEEE, 2018, pp. 3251–3256.

[9] R. Quirynen, "Numerical simulation methods for embedded optimization," dissertation, KU Leuven, 2017.

[10] S. Gros, M. Zanon, R. Quirynen, A. Bemporad, and M. Diehl, "From linear to nonlinear mpc: bridging the gap via the real-time iteration," *International Journal of Control*, vol. 93, no. 1, pp. 62–80, 2020.

[11] R. Quirynen and S. Di Cairano, "PRESAS: Block-structured preconditioning of iterative solvers within a primal active-set method for fast model predictive control," *Optimal Control Applications and Methods*.

[12] X. Feng, S. Di Cairano, and R. Quirynen, "Inexact adjoint-based sqp algorithm for real-time stochastic nonlinear mpc," in *IFAC World Congress*, 2020.

[13] M. Bujarbaruah, X. Zhang, H. E. Tseng, and F. Borrelli, "Adaptive MPC for autonomous lane keeping," *arXiv preprint arXiv:1806.04335*, 2018.

[14] F. Lin, Y. Chen, Y. Zhao, and S. Wang, "Path tracking of autonomous vehicle based on adaptive model predictive control," *International Journal of Advanced Robotic Systems*, vol. 16, no. 5, p. 1729881419880089, 2019.

[15] B.-C. Chen, B.-C. Luan, and K. Lee, "Design of lane keeping system using adaptive model predictive control," in *2014 IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE, 2014, pp. 922–926.

[16] K. Berntorp and S. Di Cairano, "Tire-stiffness and vehicle-state estimation based on noise-adaptive particle filtering," *IEEE Transactions on Control Systems Technology*, vol. 27, no. 3, pp. 1100–1114, 2018.

[17] K. Berntorp, B. Olofsson, K. Lundahl, and L. Nielsen, "Models and methodology for optimal trajectory generation in safety-critical road–vehicle manoeuvres," *Vehicle System Dynamics*, vol. 52, no. 10, pp. 1304–1332, 2014.

[18] B. Pacejka Hans, "Tire and vehicle dynamics," *Warrendale: SAE*, 2006.

[19] R. Rajamani, *Vehicle dynamics and control*. Springer Science & Business Media, 2011.

[20] D. Telen, M. Vallerio, L. Cabianca, B. Houska, J. Van Impe, and F. Logist, "Approximate robust optimization of nonlinear systems under parametric uncertainty and process noise," *Journal of Process Control*, vol. 33, pp. 140–154, 2015.

[21] S. Khaleghian, A. Emami, and S. Taheri, "A technical survey on tire-road friction estimation," *Friction*, vol. 5, no. 2, pp. 123–146, 2017.

[22] ISO 3888-2: 2002, "Passenger cars–test track for a severe lane change manoeuvre–part 2: Obstacle avoidance," 2002.