Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 54-6 (2021) 50–57

Zero-Order Robust Nonlinear Model Predictive Control with Ellipsoidal Uncertainty Sets

Andrea Zanelli * **Jonathan Frey** ** **Florian Messerer** *** **Moritz Diehl** **

* *Andrea Zanelli is with the Institute for Dynamic Systems and Control, ETH Zurich, Switzerland (e-mail: zanellia@ethz.ch)*
** *Moritz Diehl and Jonathan Frey are with the Department of Microsystems Engineering (IMTEK) and the Department of Mathematics, University of Freiburg, Germany (e-mail: {jonathan.frey, moritz.diehl}@imtek.uni-freiburg.de)*
*** *Florian Messerer is with the Department of Microsystems Engineering (IMTEK), University of Freiburg, Germany (e-mail: florian.messerer@imtek.uni-freiburg.de)*

Abstract: In this paper, we propose an efficient zero-order algorithm that can be used to compute an approximate solution to robust optimal control problems (OCP) and robustified nonconvex programs in general. In particular, we focus on robustified OCPs that make use of ellipsoidal uncertainty sets and show that, with the proposed zero-order method, we can efficiently obtain suboptimal, but robustly feasible solutions. The main idea lies in leveraging an inexact sequential quadratic programming (SQP) algorithm in which an advantageous sparsity structure is enforced. The obtained sparsity allows one to eliminate the variables associated with the propagation of the ellipsoidal uncertainty sets and to solve a reduced problem with the same dimensionality and sparsity structure of a nominal OCP. The inexact algorithm can drastically reduce the computational complexity of the SQP iterations (e.g., in the case where a structure exploiting interior-point method is used to solve the underlying quadratic programs (QPs), from $O(N \cdot (n_x^6 + n_u^3))$ to $O(N \cdot (n_x^3 + n_u^3)))$. Moreover, standard embedded QP solvers for nominal problems can be leveraged to solve the reduced QP.

Keywords: model predictive control, numerical optimization, robust optimization

1. INTRODUCTION

Although the closed-loop application of nominal nonlinear model predictive control (NMPC) allows it to react to model-plant mismatch, uncertainty cannot be taken into account explicitly in the prediction. This is especially problematic in the context of safety-critical applications, where there can be constraints that should be violated under no circumstance. Robust model predictive control (RMPC) includes the uncertainty in the prediction model and – under the assumption of finite uncertainty support – finds a control trajectory such that the constraints are fulfilled for any realization of the disturbances (see, e.g., (Morari, 1987; Rawlings et al., 2017)). Among other approaches, RMPC can be implemented using the so-called tube-based approach, in which a tube of trajectories is controlled instead of only the nominal trajectory (Mayne et al., 2011, 2005). Ideally this tube contains exactly the set of all possible trajectories and not more, but in practice it needs to be approximated, e.g, by ellipsoidal sets around the nominal trajectory (Houska, 2011). Knowledge of this tube can then be used to formulate a nominal NMPC problem with tightened constraints to account for the uncertainty. The tightening procedure is in general expensive and needs to be performed off-line, though online tightening would allow the solution to be less conservative (Köhler et al., 2018). An alternative approach would be to assume an ellipsoidal representation of the uncertainty sets over the entire prediction horizon and include the dynamics that describe their propagation – sometimes referred to as Lyapunov dynamics – into the MPC problem. The resulting ellipsoids can be defined by a symmetric and positive semidefinite matrix $\Sigma \in \mathbb{S}_+^{n_x \times n_x}$, with dynamics derived from the prediction model (Gillis and Diehl, 2013; Houska, 2011). If these dynamics are propagated alongside the nominal dynamics, the dimension of the augmented state is quadratic in n_x. This is often incompatible with an efficient online solution of the MPC problems. Furthermore, the ellipsoid dynamics are usually defined using the sensitivities of the prediction model, such that computation of their Jacobian requires second-order derivatives of the prediction model, and computation of their Hessian even requires third-order derivatives.

⋆ This research was supported by the German Federal Ministry for Economic Affairs and Energy (BMWi) via DyConPV (0324166B), by DFG and via Research Unit FOR 2401 and project 424107692 and by the EU via ITN-AWESCO (642 682).

10.1016/j.ifacol.2021.08.523

1.1 Contribution and related work

Following the ideas proposed in (Zanelli et al., 2016, 2019a), we present a zero-order approach, that we call the zero-order robust optimization (zoRO) algorithm. The strategy neglects the sensitivities of the ellipsoid dynamics in order to reduce the computational complexity of the robustified NMPC problem to roughly the same complexity of solving a nominal NMPC problem. This comes at the cost of converging to a suboptimal, but feasible, solution of the robustified NMPC problem. We corroborate our findings on a nontrivial benchmark where large speedups can be achieved with respect to a standard implementation of RMPC with ellipsoidal uncertainty sets.

A similar approach to the one presented in this section, was recently proposed in (Feng et al., 2020). With respect to (Feng et al., 2020), we additionally present a convergence proof for the zero-order iterates (which holds even at points where strict complementarity does not hold) and study the asymptotic behavior of the suboptimal solution as the uncertainty shrinks to zero. In particular, we show that the suboptimal solution deviates from the optimal one as $O(\sigma)$ with respect to the amount of uncertainty σ. Finally, although in a stochastic rather than robust setting, a related strategy is adopted in (Hewing et al., 2020, Section V), where, at every sampling time, fixed back-offs are computed by propagation of the Lyapunov dynamics for fixed state and input trajectories associated with the solution obtained at the previous sampling time. In the present work, we update the back-offs in a similar fashion, but across iterations of the optimizer, in order to recover feasibility with respect to the robustified constraints.

1.2 Notation

Throughout the paper we will denote the Euclidean norm by $\|\cdot\|$, when referring to vectors, and, with the same notation, to the spectral norm

$$\|A\| := \sqrt{\lambda_{\max}(A^\top A)}, \tag{1}$$

when referring to a (real) matrix A. All vectors are column vectors and we denote the concatenation of two vectors by

$$(x, y) := \begin{bmatrix} x \\ y \end{bmatrix}. \tag{2}$$

We denote the derivative (gradient) of any function by $\nabla f(x) = \frac{\partial f}{\partial x}(x)^\top$ and the Euclidean ball of radius r centered at x as $\mathcal{B}(x, r) := \{y : \|x - y\| \leq r\}$. For a matrix $Q \in \mathbb{R}^{n \times m}$, we denote by $q = \mathrm{vec}(Q)$, with $q \in \mathbb{R}^{n \cdot m}$, the vectorization of the matrix Q, i.e., the vector containing its stacked columns. For some positive integer n, we denote by $\mathbb{S}^n_{++}$ ($\mathbb{S}^n_{+}$) the set of symmetric positive definite (semidefinite) matrices of dimension n. For any $Q \in \mathbb{S}^n_{+}$ and any vector $q \in \mathbb{R}^n$, we denote by $\mathcal{E}(Q, q)$ the ellipsoid defined as

$$\mathcal{E}(Q, q) := \{q + Q^{\frac{1}{2}} v \mid \exists v \in \mathbb{R}^n : v^\top v \leq 1\}. \tag{3}$$

We sometimes use the shorthand $\mathcal{E}(Q) := \mathcal{E}(Q, 0)$. We denote the Minkowski sum of two sets by $\oplus$, i.e. $A \oplus B = \{a + b : a \in A, b \in B\}$. Finally, we denote the identity matrix by $\mathbb{I}$.

2. ROBUST OPTIMAL CONTROL WITH ELLIPSOIDAL UNCERTAINTY SETS

We will be concerned with the problem of controlling the discrete-time system

$$x_+ = \psi(x, u, w), \tag{4}$$

where $x \in \mathbb{R}^{n_x}$ and $u \in \mathbb{R}^{n_u}$ represent the states and inputs of the system, respectively, and $w \in \mathcal{E}(W)$, with $W \in \mathbb{S}^{n_w}_{++}$, is an uncertain disturbance. We assume that $\psi : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_w}$ is three times continuously differentiable. Moreover, we assume that the initial state of the system is contained in an ellipsoid, i.e., $x_{\mathrm{true}} \in \mathcal{E}(\bar{\Sigma}, x)$, with $\bar{\Sigma} \in \mathbb{S}^{n_x}_{++}$, centered at x.

Since constructing computationally tractable robust positive invariant tubes for nonlinear systems is in general a challenging task, we leverage the theory developed in (Houska, 2011) in order to approximate the true reachable sets with the ellipsoids obtained with a Lyapunov difference equation.

To this end, we regard the following robustified optimal control problem with ellipsoidal uncertainty sets, which is often used as an approximation to robust (Houska, 2011) and stochastic (Gillis and Diehl, 2013; Hewing et al., 2020) MPC problems:

$$\begin{aligned}
\min_{\substack{s_0,\ldots,s_N \\ \Sigma_0,\ldots,\Sigma_N \\ u_0,\ldots,u_{N-1}}} \quad & \sum_{i=0}^{N-1} l(s_i, u_i) + m(s_N) \\
\text{s.t.} \quad & s_0 - x = 0, \\
& \Sigma_0 - \bar{\Sigma} = 0, \\
& \psi(s_i, u_i, 0) - s_{i+1} = 0, \qquad i = 0, \ldots, N-1, \\
& \Sigma_{i+1} = \Phi(\Sigma_i, W_i, s_i, u_i), \quad i = 0, \ldots, N-1, \\
& \pi_i(s_i, u_i, \Sigma_i) \leq 0, \qquad i = 0, \ldots, N-1, \\
& \pi_N(s_N, \Sigma_N) \leq 0,
\end{aligned} \tag{5}$$

where s and u describe the nominal state and input trajectory, respectively. The functions $l : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}$ and $m : \mathbb{R}^{n_x} \to \mathbb{R}$ denote the cost terms. We have introduced the robustified constraints $\pi_{i,j}$ based on the nominal constraints $\hat{\pi}_{i,j}$, for $i = 0, \ldots, N-1$, with components, for $j = 1, \ldots, n_\pi$,

$$\begin{aligned}
\pi_{i,j}(s_i, u_i, \Sigma_i) := {} & \hat{\pi}_{i,j}(s_i, u_i) \\
& + \sqrt{\nabla_s \hat{\pi}_{i,j}(s_i, u_i)^\top \Sigma_i \nabla_s \hat{\pi}_{i,j}(s_i, u_i)}
\end{aligned} \tag{6}$$

and, for $j = 1, \ldots, n_{\pi_N}$,

$$\begin{aligned}
\pi_{N,j}(s_N, \Sigma_N) := {} & \hat{\pi}_{N,j}(s_N) \\
& + \sqrt{\nabla_s \hat{\pi}_{N,j}(s_N)^\top \Sigma_N \nabla_s \hat{\pi}_{N,j}(s_N)}.
\end{aligned} \tag{7}$$

The matrices Σ_i, for $i = 0, \ldots, N$, describe the uncertainty sets obtained through the ellipsoidal dynamics defined by the function $\Phi : \mathbb{R}^{n_x \times n_x} \times \mathbb{R}^{n_w \times n_w} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x \times n_x}$ defined as follows:

$$\begin{aligned}
\Phi(\Sigma, W, s, u) := {} & \\
& C(s, u)\, \Sigma\, C(s, u)^\top + B(s, u)\, W\, B(s, u)^\top.
\end{aligned} \tag{8}$$

Here we have introduced the Jacobians of the dynamics

$$C(s, u) := \frac{\partial \psi}{\partial s}(s, u, 0) \quad \text{and} \quad B(s, u) := \frac{\partial \psi}{\partial w}(s, u, 0).$$

We will refer to the square root terms in (2) and (6) as "back-off" terms. Finally, $x \in \mathbb{R}^{n_x}$ is a parameter that represents the initial state of the system.

2.1 Approximate robustness

It is well known that the reachable set of linear and nonlinear systems with ellipsoidal uncertainty is not an ellipsoid (see, e.g., Houska (2011)). However, problem (5) provides a reasonable trade-off between accuracy and computational effort. In particular, the Lyapunov difference equation (8) would provide the true reachable set for the linearized dynamics under the assumption that the uncertainty satisfies an $\mathcal{L}_2$ bound (cf. (Houska, 2011, Assumption 5.1)). Alternatively, we can compute an ellipsoidal outer approximation to the true reachable set by exploiting the fact (see, e.g., (Houska, 2011, Theorem 2.4)) that, for any $\alpha_1, \alpha_2 > 0$ such that $\alpha_1 + \alpha_2 \leq 1$ and any $W_1, W_2 \in \mathbb{S}_+^n$ it holds that $\mathcal{E}(W_1) \oplus \mathcal{E}(W_1) \subseteq \mathcal{E}(\alpha_1^{-1}W_1 + \alpha_2^{-1}W_2)$. Exploiting this fact, we can modify (8) into

$$\Phi(\Sigma, W, s, u) := \alpha_1^{-1} C(s,u)\, \Sigma\, C(s,u)^\top + \alpha_2^{-1} B(s,u)\, W\, B(s,u)^\top.$$

This tailored Lyapunov difference equation (for which we could easily modify the algorithm proposed in this paper) would provide an outer approximation to the reachable set associated with the linearized dynamics.

Similarly, following (Houska, 2011, Lemma 5.1), we can construct an outer approximation of the reachable sets (again to be interpreted in a "linearized" sense), which requires the integration of a Lyapunov ordinary differential equation (ODE). Choosing $\Omega_\tau(\lambda) = \mathcal{E}(W)$, $\mu = 1$ in (Houska, 2011, Lemma 5.1) we can interpret (8) as a discretization of the resulting Lyapunov ODE. Yet another useful interpretation of (8) is the one associated with an approximate propagation of a Gaussian distribution based on the linearized dynamics. This is the point of view taken in (Hewing et al., 2020) and (Feng et al., 2020) and it leads to an interpretation of the robustified constraints as single chance constraints.

Finally, it is worth noticing that exact robust tubes can be computed by integrating the Lyapunov ODE in (Houska, 2011, Theorem 5.3), which requires however lumping the nonlinearity in the dynamics into an ellipsoidal term, which is less practical from a computational point of view.

In all the above approaches, the approximate ellipsoidal robust tubes are used to compute a "back-off" terms in (2) and (6). However, the support of the ellipsoids is computed in the direction of the gradient of the constraints, such that the robustification (even in the case where true outer approximations to the reachable sets would be available) is again to be intended in a "linearized" sense.

3. THE ZERO-ORDER ALGORITHM FOR ROBUST OPTIMIZATION

Problem (5) can be solved with standard nonlinear programming techniques which can be substantially faster and easier to implement than the algorithms needed to solve *min-max* formulations. However, the number of optimization variables is far larger than the one associated with the nominal problem. In particular, notice that using a naive state augmentation, the number of states involved would grow from n_x to $n_x + \frac{n_x(n_x+1)}{2}$, in the case where we can exploit symmetry of the uncertainty matrices and $n_x + n_x^2$ otherwise.

In the following, we outline the proposed zero-order algorithm and show that it can efficiently obtain suboptimal, but feasible solutions to (5). The main idea lies in exploiting an approximation of the robustified OCP, where the first-order derivatives of Φ with respect to states and control are set to zero in order to enforce an advantageous sparsity structure. In this way, the obtained sparsity allows one to eliminate the variables associated with the propagation of the ellipsoidal uncertainty sets and to solve a reduced problem with the same dimensionality and sparsity structure of a nominal OCP.

In order to outline the algorithm and its properties, we utilize the vectorized uncertainty $P = \mathrm{vec}(\mathcal{P})$, with $P \in \mathbb{R}^{n_v^2}$ and $\mathcal{P} \in \mathbb{R}^{n_v \times n_v}$, and reformulate (5) into the compact formulation

$$\begin{aligned} \min_{y,P} \quad & f(y) \\ \text{s.t.} \quad & g(y) = 0, \\ & h(y) \leq 0, \\ & A(y)P - \sigma^2 \cdot b(y) = 0, \\ & \hat{h}_k(y,0) + \hat{h}_{\mathrm{back},k}(y,P) \leq 0, \quad k = 1, \dots, n_{\hat{h}}, \end{aligned} \tag{9}$$

where $y \in \mathbb{R}^{n_y}$, $f : \mathbb{R}^{n_y} \to \mathbb{R}$ and $g : \mathbb{R}^{n_y} \to \mathbb{R}^{n_g}$ and $h : \mathbb{R}^{n_y} \to \mathbb{R}^{n_h}$ are non-robustified constraints. The functions $\hat{h} : \mathbb{R}^{n_y} \times \mathbb{R}^{n_v} \to \mathbb{R}^{n_{\hat{h}}}$ and the back-off terms $\hat{h}_{\mathrm{back},k}(y,P) := \sqrt{\nabla_v \hat{h}_k(y,0)^\top \mathcal{P} \nabla_v \hat{h}_k(y,0)}$, for $k = 1, \dots, n_{\hat{h}}$, define instead the robustified constraints. We have introduced the functions $A : \mathbb{R}^{n_y} \to \mathbb{R}^{n_v^2 \times n_v^2}$ and $b : \mathbb{R}^{n_y} \to \mathbb{R}^{n_v^2}$. Finally, $\sigma \in \mathbb{R}_{++}$ is a parameter that scales the amount of uncertainty (loosely speaking, it could be interpreted as a standard deviation if the matrix $\mathcal{P}$ actually represented a covariance matrix). For simplicity of notation, we will sometimes refer to the vectorized forms

$$\hat{h}(y,0) := \begin{bmatrix} \hat{h}_1(y,0) \\ \cdots \\ \hat{h}_{n_{\hat{h}}}(y,0) \end{bmatrix}, \ \hat{h}_{\mathrm{back}}(y,P) := \begin{bmatrix} \hat{h}_{\mathrm{back},1}(y,P) \\ \cdots \\ \hat{h}_{\mathrm{back},n_{\hat{h}}}(y,P) \end{bmatrix}.$$

Problem (5) can be reformulated into (9) by observing that, e.g., for $N = 1$, the uncertainty dynamics can be rewritten as

$$\begin{bmatrix} \mathbb{I} & \mathbb{0} \\ -\Gamma_0(s_0,u_0) & \mathbb{I} \end{bmatrix} \begin{bmatrix} \Sigma_0 \\ \Sigma_1 \end{bmatrix} = \begin{bmatrix} \bar{\Sigma} \\ B(s_0,u_0)W_0B(s_0,u_0)^\top \end{bmatrix}. \tag{10}$$

Here, for any s_0, u_0, $\Gamma_0(s_0,u_0) \in \mathbb{R}^{n_x \times n_x}$ is the matrix representation of the linear map $X \to C(s_0,u_0)XC(s_0,u_0)^\top$, i.e.,

$$\Gamma_0(s_0,u_0)\Sigma_0 = C(s_0,u_0)\Sigma_0 C(s_0,u_0)^\top. \tag{11}$$

Remark 1. Notice that the zoRO algorithm introduced in the following can be applied to any robustified program of the form in (9) that need not be associated with a robustified OCP of the form in (5).

Referring to (9) will both simplify the notation and make it possible to apply the results to a more general class of robust optimization problems. We are interested in solving (9) with sequential quadratic programming (SQP). The QP subproblems to be solved at every iteration take the form

$$\begin{aligned}\min_{\Delta y, \Delta P} \quad & \nabla_y f(\hat{y})^\top \Delta y + \frac{1}{2}\begin{bmatrix}\Delta y\\ \Delta P\end{bmatrix}^\top \begin{bmatrix}M_{yy} & M_{Py}^\top\\ M_{Py} & M_{PP}\end{bmatrix}\begin{bmatrix}\Delta y\\ \Delta P\end{bmatrix}\\ \text{s.t.} \quad & g_{\text{lin}}(\hat{y}, \Delta y) = 0,\\ & h_{\text{lin}}(\hat{y}, \Delta y) \leq 0,\\ & \Phi_{\text{lin}}(\sigma;\, \hat{y}, \hat{P}, \Delta y, \Delta P) = 0,\\ & \hat{h}_{\text{lin}}(\hat{y}, \hat{P}, \Delta y, \Delta P),\end{aligned} \tag{12}$$

where

$$\begin{aligned}g_{\text{lin}}(\hat{y}, \Delta y) &:= g(\hat{y}) + \nabla g(\hat{y})^\top \Delta y,\\ h_{\text{lin}}(\hat{y}, \Delta y) &:= h(\hat{y}) + \nabla h(\hat{y})^\top \Delta y,\\ \Phi_{\text{lin}}(\sigma;\, \hat{y}, \hat{P}, \Delta y, \Delta P) &:= A(\hat{y})\hat{P} - \sigma^2 \cdot b(\hat{y})\\ &\quad + \frac{\partial}{\partial y}\left(A(\hat{y})\hat{P}\right)\Delta y + A(\hat{y})\Delta P\\ &\quad - \sigma^2 \cdot \nabla b(\hat{y})^\top \Delta y\end{aligned} \tag{13}$$

and

$$\begin{aligned}\hat{h}_{\text{lin}}(\hat{y}, \hat{P}, \Delta y, \Delta P) &:= \hat{h}(\hat{y}, 0) + \hat{h}_{\text{back}}(\hat{y}, \hat{P})\\ &\quad + \nabla_y \hat{h}(\hat{y}, 0)^\top \Delta y + \nabla_y \hat{h}_{\text{back}}(\hat{y}, \hat{P})^\top \Delta y\\ &\quad + \nabla_P \hat{h}_{\text{back}}(\hat{y}, \hat{P})^\top \Delta P.\end{aligned} \tag{14}$$

Finally,

$$M = \begin{bmatrix}M_{yy} & M_{Py}^\top\\ M_{Py} & M_{PP}\end{bmatrix} \tag{15}$$

represents the chosen symmetric and positive definite Hessian approximation. Due to the introduction of the vectorized uncertainty matrix P, solving the subproblems in (12) can be substantially more computationally demanding than solving their nominal counterpart. For this reason, we propose to instead solve modified subproblems of the form

$$\begin{aligned}\min_{\Delta y, \Delta P} \quad & \nabla f(\hat{y})^\top \Delta y + \frac{1}{2}\begin{bmatrix}\Delta y\\ \Delta P\end{bmatrix}^\top \begin{bmatrix}M_{yy} & M_{Py}^\top\\ M_{Py} & M_{PP}\end{bmatrix}\begin{bmatrix}\Delta y\\ \Delta P\end{bmatrix}\\ \text{s.t.} \quad & g_{\text{lin}}(\hat{y}, \Delta y) = 0,\\ & h_{\text{lin}}(\hat{y}, \Delta y) \leq 0,\\ & \tilde{\Phi}_{\text{lin}}(\sigma;\, \hat{y}, \hat{P}, \Delta P) = 0,\\ & \hat{h}_{\text{lin}}(\hat{y}, \hat{P}, \Delta y, \Delta P),\end{aligned} \tag{16}$$

where Φ_{lin} has been replaced by

$$\tilde{\Phi}_{\text{lin}}(\hat{y}, \hat{P}, \Delta P) := A(\hat{y})\hat{P} - \sigma^2 \cdot b(\hat{y}) + A(\hat{y})\Delta P. \tag{17}$$

Remark 2. The approximation of the first-order derivatives in (17) is in line with the zero-order strategies analyzed in (Bock et al., 2007; Zanelli et al., 2016, 2019b). Notice however, that in the zoRO algorithm we do evaluate exactly part of the derivatives.

Under the assumption that $A(\hat{y})$ is invertible, we can then eliminate ΔP as

$$\Delta\tilde{P} := -\hat{P} + A(\hat{y})^{-1}\sigma^2 \cdot b(\hat{y}) \tag{18}$$

such that the QP subproblems take the form

$$\begin{aligned}\min_{\Delta y} \quad & (\nabla f(\hat{y})^\top + \Delta\tilde{P}^\top M_{Py})\Delta y + \frac{1}{2}\Delta y^\top M_{yy}\Delta y\\ \text{s.t.} \quad & g_{\text{lin}}(\hat{y}, \Delta y) = 0,\\ & h_{\text{lin}}(\hat{y}, \Delta y) \leq 0,\\ & \hat{h}_{\text{lin}}(\hat{y}, \hat{P}, \Delta y, \Delta\tilde{P}) \leq 0.\end{aligned} \tag{19}$$

The elimination of ΔP described in (18) together with the solution of the reduced problem (18) can be considerably less computationally demanding than solving the original problem (9). In particular, in the context of NMPC, ΔP can be eliminated through the recursion

$$\Delta\Sigma_{i+1} = \Gamma_i(\hat{s}_i, \hat{u}_i)\hat{\Sigma}_i + B_i W_i B_i^\top - \hat{\Sigma}_i, \tag{20}$$

for $i = 0, \ldots, N-1$, and the reduced problem takes the form of a nominal OCP, hence with a drastically decreased number of states.

3.1 Asymptotic analysis of the zoRO solution

In order to be able to leverage the results on strongly regular generalized equations, which, in their standard form, require Lipschitz continuity of the underlying functions, we reformulate (9) as

$$\begin{aligned}\min_{y,Q} \quad & f(y)\\ \text{s.t.} \quad & g(y) = 0,\\ & h(y) \leq 0,\\ & A(y)Q - b(y) = 0,\\ & \hat{h}(y, 0) + \sigma \cdot \hat{h}_{\text{back}}(y, Q) \leq 0,\end{aligned} \tag{21}$$

where we have used the substitution $P = \sigma^2 \cdot Q$. In this was, the back-off terms are not necessarily evaluated at 0 for $\sigma \to 0$. Let λ, μ, ν and η denote the Lagrange multipliers associated with the constraints in (16) and define

$$\begin{aligned}&\nabla_{(y,Q)}\tilde{\mathcal{L}}(\sigma;\, y, Q, \lambda, \mu, \nu, \eta) :=\\ &\quad + \begin{bmatrix}\nabla f(y)\\ 0\end{bmatrix} + \begin{bmatrix}\nabla g(y)\\ 0\end{bmatrix}\lambda\\ &\quad + \begin{bmatrix}\nabla h(y)\\ 0\end{bmatrix}\mu + \begin{bmatrix}0\\ A(y)^\top\end{bmatrix}\nu\\ &\quad + \begin{bmatrix}\nabla_y \hat{h}(y, 0) + \sigma \cdot \nabla_y \hat{h}_{\text{back}}(y, Q)\\ \sigma \cdot \nabla_Q \hat{h}_{\text{back}}(y, Q)\end{bmatrix}\eta,\end{aligned} \tag{22}$$

where, with respect to the exact gradient of the Lagrangian, we have omitted the term $\frac{\partial}{\partial y}\left(A(y)Q\right)^\top$. If the zoRO iterates achieve convergence, i.e., $\Delta y = 0$ and $\Delta P = 0$, from the first-order optimality conditions of (16) we obtain, after the substitution $P = \sigma^2 \cdot Q$, that the solution must satisfy the following system:

$$\begin{aligned}&\nabla_{(y,Q)}\tilde{\mathcal{L}}(\sigma;\, y, Q, \lambda, \mu, \nu, \eta) = 0,\\ &g(y) = 0,\\ &h(y) \leq 0,\\ &A(y)Q - b(y) = 0,\\ &\hat{h}(y, 0) + \sigma \cdot \hat{h}_{\text{back}}(y, Q) \leq 0,\end{aligned} \tag{23}$$

together with $\mu, \eta \geq 0$ and the associated complementarity constraints. The first-order optimality conditions of the original problem (9) would instead read

$$\begin{aligned}&\nabla_{(y,Q)}\mathcal{L}(\sigma;\, y, Q, \lambda, \mu, \nu, \eta) = 0,\\ &g(y) = 0,\\ &h(y) \leq 0,\\ &A(y)Q - b(y) = 0,\\ &\hat{h}(y, 0) + \sigma \cdot \hat{h}_{\text{back}}(y, Q) \leq 0,\end{aligned} \tag{24}$$

with

$$\begin{aligned}&\nabla\mathcal{L}_{(y,Q)}(\sigma;\, y, Q, \lambda, \mu, \nu, \eta) =\\ &\quad \nabla\tilde{\mathcal{L}}(\sigma;\, y, Q, \lambda, \mu, \nu, \eta) + \begin{bmatrix}\frac{\partial}{\partial y}\left(A(y)Q\right)^\top - \nabla b(y)\\ 0\end{bmatrix}\nu.\end{aligned}$$

Remark 3. Notice that, as observed in (Feng et al., 2020), it is in principle possible to compensate the deviation between the exact and approximate gradients of the Lagrangian using "adjoint" corrections as initially proposed in (Bock et al., 2007). In this fashion it is possible to obtain an algorithm that converges to locally optimal solutions. In this work, we are however interested in the properties of the approximate solution recovered by the zoRO algorithm without such corrections.

Comparing (23) and (24), we see that the solution recovered by the zoRO iterates satisfies a perturbed version of (24) which however enforces feasibility with respect to the original constraints.

In order to study the asymptotic behavior of the solution to (23) as $\sigma \to 0$, we will reformulate the above first-order optimality conditions as generalized equations. To this end, let $z := (y, Q, \lambda, \mu, \nu, \eta)$ represent the primal-dual variables and define $K := \mathbb{R}^{n_y+n_Q} \times \mathbb{R}^{n_g} \times \mathbb{R}_+^{n_h} \times \mathbb{R}^{n_v^2} \times \mathbb{R}_+^{n_{\hat{h}}}$. Following the formulation from (Robinson, 1980), we can write (24) as

$$0 \in F(\sigma;\, z) + \mathcal{N}_K(z), \tag{25}$$

where we have introduced

$$F(\sigma;\, z) := \begin{bmatrix} \nabla_{(y,Q)} \mathcal{L}(\sigma;\, y, Q, \lambda, \mu, \nu, \eta) \\ -g(y) \\ -h(y) \\ -A(y)Q + b(y) \\ -\hat{h}(y, 0) - \sigma \cdot \hat{h}_{\text{back}}(y, Q) \end{bmatrix}. \tag{26}$$

Here $\mathcal{N}_K(z)$ denotes the normal cone to the set K at z. It is easy to verify that a solution recovered by the zoRO algorithm, which we will refer to as $\tilde{z}(\sigma)$, satisfies the following inclusion

$$0 \in F(\sigma;\, z) + \epsilon(z) + \mathcal{N}_K(z) \tag{27}$$

with

$$\epsilon(z) := \begin{bmatrix} -\frac{\partial}{\partial y} \left(A(y)Q\right)^\top + \nabla b(y) \\ 0 \end{bmatrix} \nu. \tag{28}$$

Assumption 4. Let $\bar{z} : \mathbb{R} \to \mathbb{R}^{n_z};\, \sigma \to \bar{z}(\sigma)$ be a single-valued localization of the solution map of (25). Assume that (25) is strongly regular at $\bar{z}(0)$ with Lipschitz constant γ.

Theorem 5. Let Assumption 4 hold. Then there exists a single-valued and Lipschitz localization $\tilde{z} : \mathbb{R} \to \mathbb{R}^{n_z};\, \sigma \to \tilde{z}(\sigma)$ of the solution map of (27) that satisfies

$$\|\tilde{z}(\sigma) - \bar{z}(\sigma)\| = O(\sigma). \tag{29}$$

Proof. Regard the auxiliary generalized equation

$$0 \in F(\sigma;\, z) + \epsilon_{\text{p}} + \mathcal{N}_K(z). \tag{30}$$

Due to Assumption 4, there exists a locally unique single-valued and Lipschitz continuous localization $\bar{z}_{\text{a}} : \mathbb{R} \times \mathbb{R}^{n_z} \to \mathbb{R}^{n_z};\, (\sigma, \epsilon_{\text{p}}) \to \bar{z}_{\text{a}}(\sigma, \epsilon_{\text{p}})$ of its solution map. Notice that a solution $\tilde{z}(\sigma)$ to (27) is also a solution to

$$0 \in F(\sigma;\, z) + \epsilon(\tilde{z}(\sigma)) + \mathcal{N}_K(z), \tag{31}$$

such that, due to local uniqueness, for sufficiently small perturbations, $\bar{z}_{\text{a}}(\sigma, \epsilon(\tilde{z}(\sigma))) = \tilde{z}(\sigma)$ and that $\bar{z}_{\text{a}}(\sigma, 0) = \bar{z}(\sigma)$. Moreover, due to strong regularity, we have

$$\|\bar{z}_{\text{a}}(\sigma, \epsilon_{\text{p}}) - \bar{z}_{\text{a}}(0, 0)\| \leq \gamma \cdot \left\| \begin{matrix} \sigma \\ \epsilon_{\text{p}} \end{matrix} \right\| \tag{32}$$

and

$$\|\bar{z}_{\text{a}}(\sigma, 0) - \bar{z}_{\text{a}}(0, 0)\| \leq \gamma \cdot \|\sigma\|. \tag{33}$$

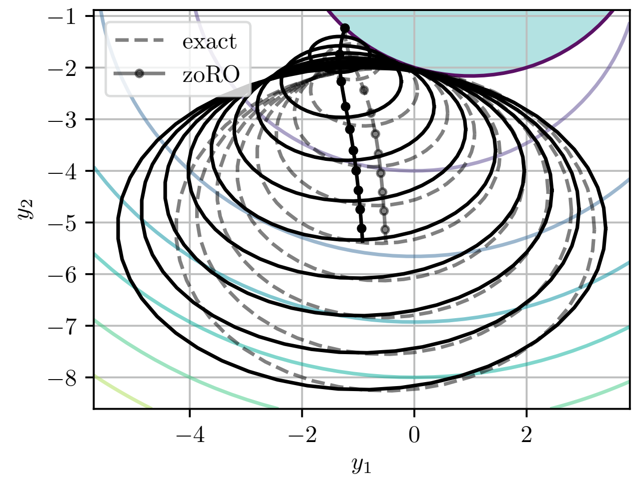

Fig. 1. Illustrative example. The complement of the feasible region is coloured in green and the level lines of the objective are reported. The solution the zoRO iterates converge to is robustly feasible and the deviation from the exact solution to (43) is proportional to σ.

This implies

$$\|\bar{z}_{\text{a}}(\sigma, \epsilon_{\text{p}}) - \bar{z}_{\text{a}}(\sigma, 0)\| \leq \gamma \cdot \left(\left\| \begin{matrix} \sigma \\ \epsilon_{\text{p}} \end{matrix} \right\| + \|\sigma\| \right). \tag{34}$$

Exploiting the easily verifiable fact that $\tilde{\nu}(0) = 0$ and that, due to strong regularity we have that $\|\tilde{\nu}(\sigma)\| = O(\sigma)$, we can write $\|\epsilon(\tilde{z}(\sigma))\| = O(\sigma)$. Finally, we obtain

$$\begin{aligned} \|\tilde{z}(\sigma) - \bar{z}(\sigma)\| &= \|\bar{z}_{\text{a}}(\sigma, \epsilon(\tilde{z}(\sigma))) - \bar{z}_{\text{a}}(\sigma, 0)\| \\ &\leq \gamma \cdot \left(\left\| \begin{matrix} \sigma \\ \epsilon(\tilde{z}(\sigma)) \end{matrix} \right\| + \|\sigma\| \right) = O(\sigma), \end{aligned} \tag{35}$$

which concludes the proof. □

Theorem 5 shows that, if convergence is achieved, the asymptotic deviation of the primal-dual solution $\tilde{z}(\sigma)$ recovered by the zoRO iterates from the optimal solution $\bar{z}(\sigma)$ is of order $O(\sigma)$. In the following, we show that under suitable assumptions the zoRO iterates actually converge to $\tilde{z}(\sigma)$.

3.2 Convergence of the zoRO iterates

In order to analyze the convergence properties of the proposed numerical strategy, we observe that the inexact SQP iterates can be interpreted as the ones generated by a generalized Newton-type method applied to a specific generalized equation. To this end, we will refer to the iterates in the original space: since we are interested in the convergence for any strictly positive σ, non-Lipschitz behaviour of the back-off term will not affect our analysis. Let $z = (y, P, \lambda, \mu, \nu, \eta)$ and

$$\tilde{F}(\sigma;\, z) := \begin{bmatrix} \nabla_{(y,P)} \tilde{\mathcal{L}}(\sigma;\, y, P, \lambda, \mu, \nu, \eta) \\ -g(y) \\ -h(y) \\ -A(y)P + \sigma^2 \cdot b(y) \\ -\hat{h}(y, 0) - \hat{h}_{\text{back}}(y, P) \end{bmatrix}. \tag{36}$$

We regard the generalized equation

$$0 \in \tilde{F}(\sigma;\, z) + \mathcal{N}_K(z), \tag{37}$$

where K is defined as in (25).

Every zoRO iteration solves the linear generalized equation

$$0 \in \tilde{F}(\sigma;\, \hat{z}) + J(\hat{z})(z - \hat{z}) + \mathcal{N}_K(z), \tag{38}$$

with linearization point $\hat{z}$ and

$$J(\hat{z}) \approx \frac{\partial \tilde{F}}{\partial z}(\sigma;\, \hat{z}). \tag{39}$$

In particular, it is easy to check that, according to the definition of the zoRO iterates, the following holds:

$$J(\hat{z}) - \frac{\partial \tilde{F}}{\partial z}(\sigma;\, \hat{z}) = \begin{bmatrix} M - \frac{\partial \tilde{F}}{\partial (y,P)}(\sigma;\, \hat{z}) & 0 \\ 0 & 0 \\ \frac{\partial}{\partial y}\left(A(\hat{y})\hat{P}\right) - \sigma^2 \nabla b(\hat{y})^\top & 0 \\ 0 & 0 \end{bmatrix}. \tag{40}$$

We make the following assumption.

Assumption 6. Let $\bar{z}$ be a solution to (37). Assume that (37) is strongly regular at $\bar{z}$ with Lipschitz constant γ over the neighborhood $\mathcal{B}(\bar{z}, \tilde{r}_z)$, with $\tilde{r}_z > 0$.

Assumption 7. Assume that there exist a non-empty neighborhood $\mathcal{B}(\bar{z}, \hat{r}_z)$, with $\hat{r}_z < \tilde{r}_z$, and a positive constant $\tilde{\kappa}$, with $\gamma\tilde{\kappa} < \frac{1}{2}$ such that, for any $\hat{z} \in \mathcal{B}(\bar{z}, \hat{r}_z)$, the following holds:

$$\left\| \frac{\partial \tilde{F}}{\partial z}(\hat{z}) - J(\hat{z}) \right\| \leq \tilde{\kappa}. \tag{41}$$

Remark 8. Under Assumption 6, for the subcomponent $\bar{P}(\sigma)$ of a solution to (37), we can write $\bar{P}(\sigma) = O(\sigma)$. Hence, if M is a "sufficiently" good approximation of $\frac{\partial \tilde{F}}{\partial (y,P)}(\sigma;\, \hat{z})$, Assumption 7 can be expected to be satisfied for σ sufficiently small.

Theorem 9. Let Assumptions 6 and 7 hold and let z_+ denote the solution to (38) constructed at the linearization point $\hat{z}$. Then, there exist a strictly positive constant r_z and a positive constant $\kappa < 1$, such that, for any $\hat{z}$ in $\mathcal{B}(\bar{z}, r_z)$, the following holds:

$$\|z_+ - \bar{z}\| \leq \kappa \|\hat{z} - \bar{z}\|. \tag{42}$$

Proof. The proof follows standard arguments for generalized Newton-type methods (see, e.g., (Zanelli et al., 2019b)).

3.3 Illustrative example

In order to illustrate the properties of the proposed zoRO algorithm, we regard the following simple robustified nonlinear program:

$$\begin{aligned} \min_{y,P} \quad & f(y) \\ \text{s.t.} \quad & A(y)P - \sigma^2 \cdot b(y) = 0, \\ & \hat{h}(y, 0) + \hat{h}_{\text{back}}(y, P) \leq 0. \end{aligned} \tag{43}$$

Here $y, v \in \mathbb{R}^2$, $f(y) := \frac{1}{2} y^\top M y$ and $\hat{h}(y, v) := \sqrt{(y_1 + v_1 - 1)^2 + (y_2 + v_2 - 1)^2}$. The function $A(y)$ and $b(y)$ have been defined as $A(y) :=$

$$\begin{bmatrix} 1 + \alpha \cdot \sin(y_1) & 0 & 0 & 0 \\ 0 & 1 + \alpha \cdot \cos(y_2) & 0 & 0 \\ 0 & 0.1 + \sin(y_2) & 1 + \alpha \cdot \sin(y_2^2) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

where $\alpha \in \mathbb{R}$ is a tuning parameter, and $b(y) := (1, 0, 0, 1)$.

The zoRO algorithm has been prototyped (code available at `https://github.com/FreyJo/zoro-NMPC-2021`) in Python using `CasADi` (Andersson et al., 2018) and its interfaces to `qpOASES` (Ferreau et al., 2014) and `Ipopt` (Wächter and Biegler, 2006). Problem (43) has been solved for different values of $\sigma \in (0, \frac{3}{2}\sqrt{5}]$ and $\alpha = 0.6$ with both the zoRO algorithm and `Ipopt` applied directly to the original formulation. Notice that the values for σ and α have been picked for satisfactory visualization.

The obtained optimal solution is reported in Figure 1, where we see that the zoRO solution is suboptimal, but still satisfies the constraints in (43). For each value of σ considered we plot a dot and the associated ellipsoid defined by P. Notice that, for the sake of clarity of visualization, we use a definition of the constraints that ensures an exact robustification, i.e., the ellipsoids are "touching" the boundaries of the feasible set. This is generally not the case as the direction along which the support of the uncertainty sets is evaluated is the one of the gradient of the constraints and only provides a linear approximation. This would however affect both the exact and the zoRO solution to (9).

4. NUMERICAL EXAMPLE

In this section, we consider the task of controlling a nonlinear hanging chain model, following (Wirsching et al., 2006; Kouzoupis et al., 2018).

The system of interest consists of n_{mass} masses connected by springs. The mass on the one end of the chain is fixed, while the velocity of the mass on the other end is the control input. The model has $n_x = 6(n_{\text{mass}} - 2) + 3$ states that correspond to the velocity and position vectors of the intermediate masses and the position of the actuated mass. The chain is moved from its rest position by applying the control input $[-1, 1, 1]$ for 1 s. The controller is required to stabilize the chain at its rest position, while respecting a constraint that describes a wall at $y = -0.05$. The system is disturbed in each time step by adding disturbances to the velocities of the masses, which are randomly sampled from the ellipsoid $\mathcal{E}(10^2 \cdot \mathbb{I})$.

All variants are implemented using the Python interface of the open-source software package `acados`, which provides high-performance SQP-type methods for NMPC applications (Verschueren et al., 2019). The QPs are solved using the high-performance QP solver `HPIPM` (Frison and Diehl, 2020) without condensing and `BLASFEO` (Frison et al., 2018) with the dedicated linear algebra target. The code is available at `https://github.com/FreyJo/zoro-NMPC-2021`.

Variants: We regard the following controller variants that are compared with respect to constraint satisfaction and computational complexity:

- **nominal**: The wall constraint is imposed without the uncertainty back-off in equation (2).
- **robust**: The wall constraint is imposed with the uncertainty back-off in equation (2). The uncertainty matrices Σ_i are propagated within the integrator by augmenting the model with the continuous-time Lyapunov dynamics

$$\dot{\Sigma}(t) = C(t)\Sigma(t) + \Sigma(t)C(t)^\top + B(t)W(t)B(t)^\top \tag{44}$$

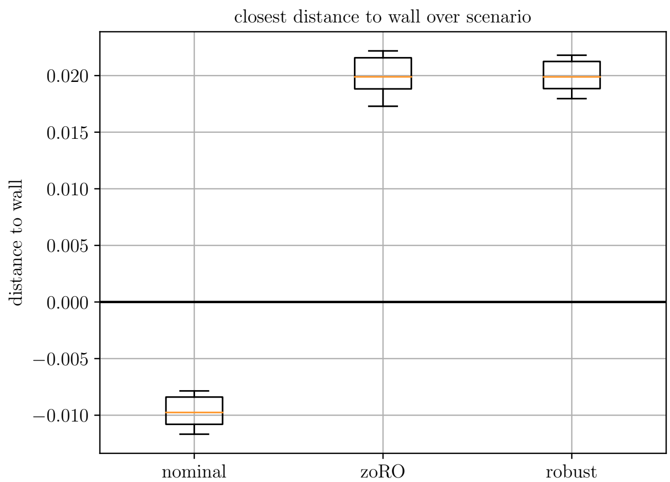

Fig. 2. Closest distance to the wall over a scenario of disturbances. Boxplot obtained over a variety of random sampled uncertainties. Converged SQP with $n_{\mathrm{mass}} = 5$

exploiting only the symmetry of Σ.

- **zoRO**: The wall constraint is imposed with the uncertainty back-off in equation (2) using the zero-order algorithm proposed in Section 3. The propagation of Σ_i is carried out in native Python in between subsequent SQP iterations.

Timings In order to validate the computational efficiency of the proposed strategy, we use the MPC controllers based on the three variants in a closed-loop simulation. In Figure 2, the minimal distance of the chain to the wall over a variety of scenarios is visualized as a boxplot. It can be observed, that while the wall constraint is violated at some point in all scenarios for the nominal controller, the two robust controllers never violate the constraint in any of the scenarios. The minor difference between the zoRO and the robust variants is due to *i)* the zero-order version converging to a slightly suboptimal solution and *ii)* the disturbances entering the simulation slightly differently. More precisely, the disturbances enter in continuous-time via the integrator for the robust version and in a discrete-time fashion in the ZOro version.

Remark 10. Notice that, since the formulation in (5) is based on an approximate propagation of the uncertainty, we cannot claim that the proposed strategy is robust with respect to disturbances drawn from a specific (bounded) set. The goal of this numerical study is instead to compare the computation times associated with different algorithms used to solve (5).

All computations were carried out on a Lenovo Carbon-X1 (7gen) Laptop with Intel i7-8565U CPU, 16 GB RAM, 1TB SSD running Ubuntu 20.04.
In Figure 3, the computation times for the different variants are visualized for $n_{\mathrm{mass}} = 7$. One can observe that the zero-order robust controller requires a slightly higher CPU time to evaluate the feedback policy compared to the nominal MPC controller, but these computation times are of a similar magnitude. In contrast to that, the exact robust controller, needs a lot more computation time, namely roughly a factor of 1000 more than the other controllers.

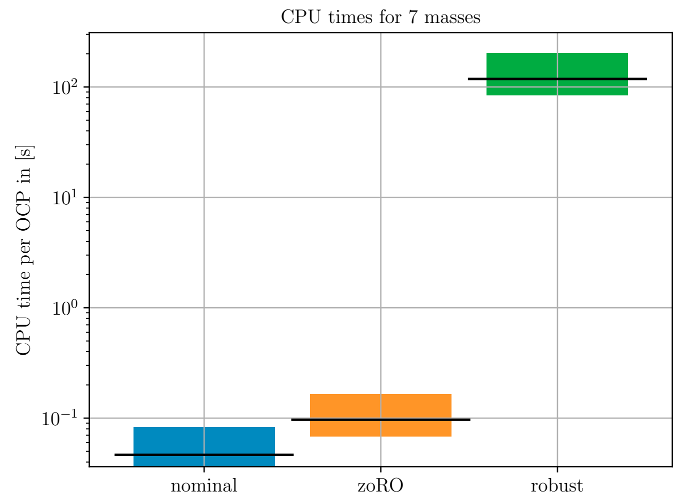

Fig. 3. CPU time per OCP. The lower and upper end show the minimum, respectively maximum CPU time over all solver calls. The black line shows the mean value of the CPU times. The proposed zero-order strategy drastically reduces the computation times and achieves, for this scenario, a speedup of about three orders of magnitude with respect to a standard implementation of robust NMPC.

Figure 4 shows how the CPU time of the controller variants scales with the number of masses n_{mass}, respectively the state dimension n_x. As an orientation, level lines corresponding to $O(n_x^3)$ and $O(n_x^6)$ are included. One can observe that the CPU time for the nominal and the zoRO controllers scales roughly as $O(n_x^3)$, while for the exact robust implementation, the computation time scales similar to $O(n_x^6)$.

Notice that, although in this context we use the feedback policy associated with "converged" solutions, i.e. obtained with a sufficiently high number of iterations, the zoRO algorithm could be in principle used in its *real-time* variant. In this case, in order to reduce the computation times, a single SQP iteration is carried out per sampling time and the approximate solution is used to control the system in the spirit of the real-time iteration (RTI) strategy (Diehl, 2001). Although this would further approximate the underlying feedback policy, under the assumptions in (Zanelli et al., 2020) it would still be possible to show asymptotic stability of the combined system-optimizer dynamics.

5. CONCLUSIONS

In this paper, we proposed an efficient inexact numerical strategy for the solution of nonconvex programs arising from robustified nonlinear model predictive control with ellipsoidal uncertainty sets. In particular, the presented strategy leverages concepts of zero-order methods that allow one to approximate first- and second-order information in the underlying numerical algorithms in order to obtain suboptimal, but feasible solutions to nonconvex programs. We showed how, in the setting of robust MPC with ellipsoidal uncertainty sets, this can be used to drastically decrease the computational complexity of the QP

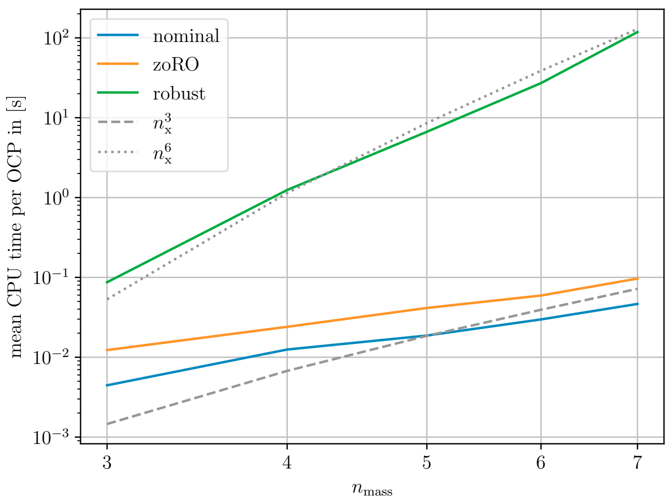

Fig. 4. Mean CPU time of an OCP solver call as a function of the number of masses. The improved asymptotic complexity of $O(n_x^3)$ is approximately matched by the measured timings.

solutions within an SQP-type algorithm. A convergence proof of the inexact iterates and a suboptimality bound for the converged solution are provided. The theoretical results were validated on a nontrivial numerical example where large speedups can be achieved with respect to a standard implementation of robust NMPC with ellipsoidal uncertainty sets.

REFERENCES

Andersson, J.A.E., Gillis, J., Horn, G., Rawlings, J.B., and Diehl, M. (2018). CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation.*

Bock, H.G., Diehl, M., Kostina, E.A., and Schlöder, J.P. (2007). Constrained optimal feedback control of systems governed by large differential algebraic equations. In *Real-Time and Online PDE-Constrained Optimization*, 3–22. SIAM.

Diehl, M. (2001). *Real-Time Optimization for Large Scale Nonlinear Processes.* Ph.D. thesis, University of Heidelberg.

Feng, X., Cairano, S.D., and Quirynen, R. (2020). Inexact Adjoint-based SQP Algorithm for Real-Time Stochastic nonlinear MPC. In *Proceedings of the IFAC World Congress.*

Ferreau, H.J., Kirches, C., Potschka, A., Bock, H.G., and Diehl, M. (2014). qpOASES: a parametric active-set algorithm for quadratic programming. *Mathematical Programming Computation*, 6(4), 327–363.

Frison, G. and Diehl, M. (2020). HPIPM: a high-performance quadratic programming framework for model predictive control. In *Proceedings of the IFAC World Congress.* Berlin, Germany.

Frison, G., Kouzoupis, D., Sartor, T., Zanelli, A., and Diehl, M. (2018). BLASFEO: Basic linear algebra subroutines for embedded optimization. *ACM Transactions on Mathematical Software (TOMS)*, 44(4), 42:1–42:30.

Gillis, J. and Diehl, M. (2013). A positive definiteness preserving discretization method for nonlinear Lyapunov differential equations. In *Proceedings of the IEEE Conference on Decision and Control (CDC).*

Hewing, L., Kabzan, J., and Zeilinger, M. (2020). Cautious model predictive control using gaussian process regression. *Transactions on Control Systems Technology*, 28(6), 2736–2743.

Houska, B. (2011). *Robust Optimization of Dynamic Systems.* Ph.D. thesis, KU Leuven. (ISBN: 978-94-6018-394-2).

Köhler, J., Müller, M.A., and Allgöwer, F. (2018). A novel constraint tightening approach for nonlinear robust model predictive control. *Annual American Control Conference (ACC).*

Kouzoupis, D., Frison, G., Zanelli, A., and Diehl, M. (2018). Recent advances in quadratic programming algorithms for nonlinear model predictive control. *Vietnam Journal of Mathematics*, 46(4), 863–882.

Mayne, D., Kerrigan, E., van Wyk, E.J., and Falugi, P. (2011). Tube-based robust nonlinear model predictive control. *International Journal of Robust and Nonlinear Control*, 21, 1341–1353.

Mayne, D.Q., Seron, M.M., and Rakovic, S.V. (2005). Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41, 219–224.

Morari, M. (1987). Robust Process Control. *Chem. Eng. Res. Des.*, 65, 462–479.

Rawlings, J.B., Mayne, D.Q., and Diehl, M.M. (2017). *Model Predictive Control: Theory, Computation, and Design.* Nob Hill, 2nd edition edition.

Robinson, S.M. (1980). Strongly Regular Generalized Equations. *Mathematics of Operations Research, Vol. 5, No. 1 (Feb., 1980), pp. 43-62*, 5, 43–62.

Verschueren, R., Frison, G., Kouzoupis, D., van Duijkeren, N., Zanelli, A., Novoselnik, B., Frey, J., Albin, T., Quirynen, R., and Diehl, M. (2019). acados: a modular open-source framework for fast embedded optimal control. *arXiv preprint.* URL `https://arxiv.org/abs/1910.13753`.

Wächter, A. and Biegler, L.T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1), 25–57.

Wirsching, L., Bock, H.G., and Diehl, M. (2006). Fast NMPC of a chain of masses connected by springs. In *Proceedings of the IEEE International Conference on Control Applications, Munich*, 591–596.

Zanelli, A., Quirynen, R., and Diehl, M. (2016). An efficient inexact NMPC scheme with stability and feasibility guarantees. In *Proceedings of 10th IFAC Symposium on Nonlinear Control Systems.* Monterey, CA, USA.

Zanelli, A., Quirynen, R., and Diehl, M. (2019a). Efficient zero-order NMPC with feasibility and stability guarantees. In *Proceedings of the European Control Conference (ECC).* Naples, Italy.

Zanelli, A., Tran-Dinh, Q., and Diehl, M. (2019b). Contraction estimates for abstract real-time algorithms for NMPC. In *Proceedings of the IEEE Conference on Decision and Control.* Nice, France.

Zanelli, A., Tran-Dinh, Q., and Diehl, M. (2020). A Lyapunov function for the combined system-optimizer dynamics in inexact model predictive control. *arXiv preprint.*