ETH zürich ELO-X

Institute for Dynamic
Systems and Control

Zero-Order Optimization for Gaussian Process-based MPC

Amon Lahr, Andrea Zanelli, Andrea Carron, Melanie N. Zeilinger

February 28, 2023, Robust MPC Workshop Freiburg

Funded by the European Union

Model Predictive Control using Gaussian Processes

[Torrente, Kaufman

linger, 2019]

Uncertainty-aware high-performance control

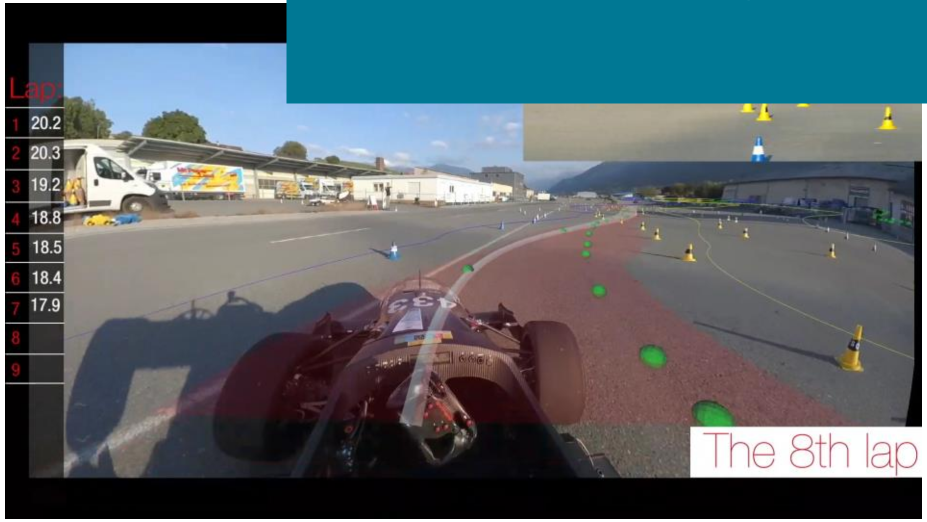

[Kabzan, Hewing, Liniger, Zeilinger, 2019]

[Ostafew, Schoellig, Barfoot, 2014]

1st Lap: Data collection with **nominal controller**

Dynamics model

Discrete-time dynamics

$$x(k+1) = \psi(x(k), u(k)) + \eta(x(k), u(k)) + w(k)$$

- $\psi(x(k), u(k))$: Nominal dynamics
- $\eta(x(k), u(k))$: Unknown dynamics
- $w(k)$: Disturbance

Dynamics model – Race car example

Discrete-time dynamics

$$x(k+1) = \psi(x(k), u(k)) + \eta(x(k), u(k)) + w(k)$$

Nonlinear bicycle model

Tire friction
Lateral forces
Aerodynamics
…

Disturbance

Dynamics model

GP model

$$d(x_i, u_i) \sim \mathcal{GP}(\mu^d, \Sigma^d)$$

Discrete-time dynamics

$$x_{i+1} = \psi(x_i, u_i) + Bd(x_i, u_i) + w_i$$

Nominal dynamics

GP model subspace

Gaussian i.i.d. disturbance

$$w_i \sim \mathcal{N}(0, \Sigma^w)$$

GP inference

Training data = Full state measurements, control inputs

$$\underbrace{B^\dagger(x_{i+1} - \psi(x_i, u_i))}_{y_{\mathrm{train},i}} = d(x_i, u_i) + B^\dagger w_i$$

GP posterior

- Non-parametric regression
- Uncertainty quantification using posterior covariance
- Cubic scaling of computational cost w.r.t. data points, $\mathcal{O}(D^3)$

→GP approximations for online evaluation (Quiñonero-Candela, Rasmussen, 2005, Snelson, Ghahramani, 2006, ..., Berntorp, 2021, Gruner et al., 2022, Vaskov et al., 2022)

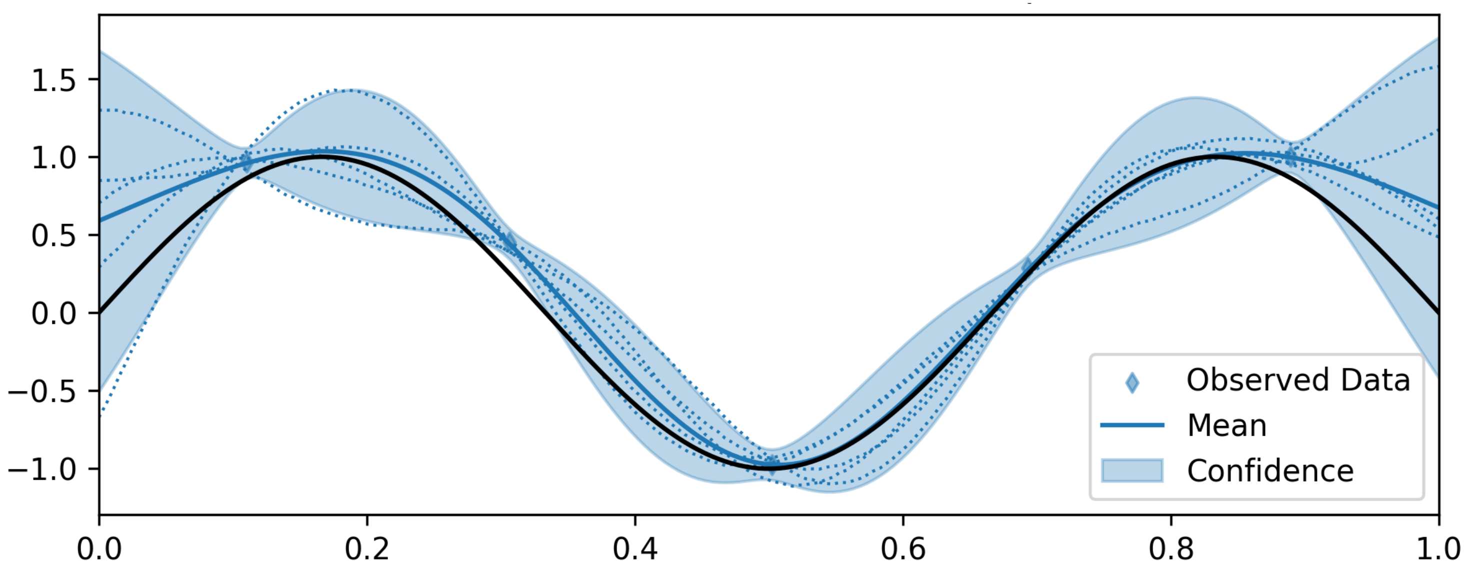

GP-MPC – Optimization problem

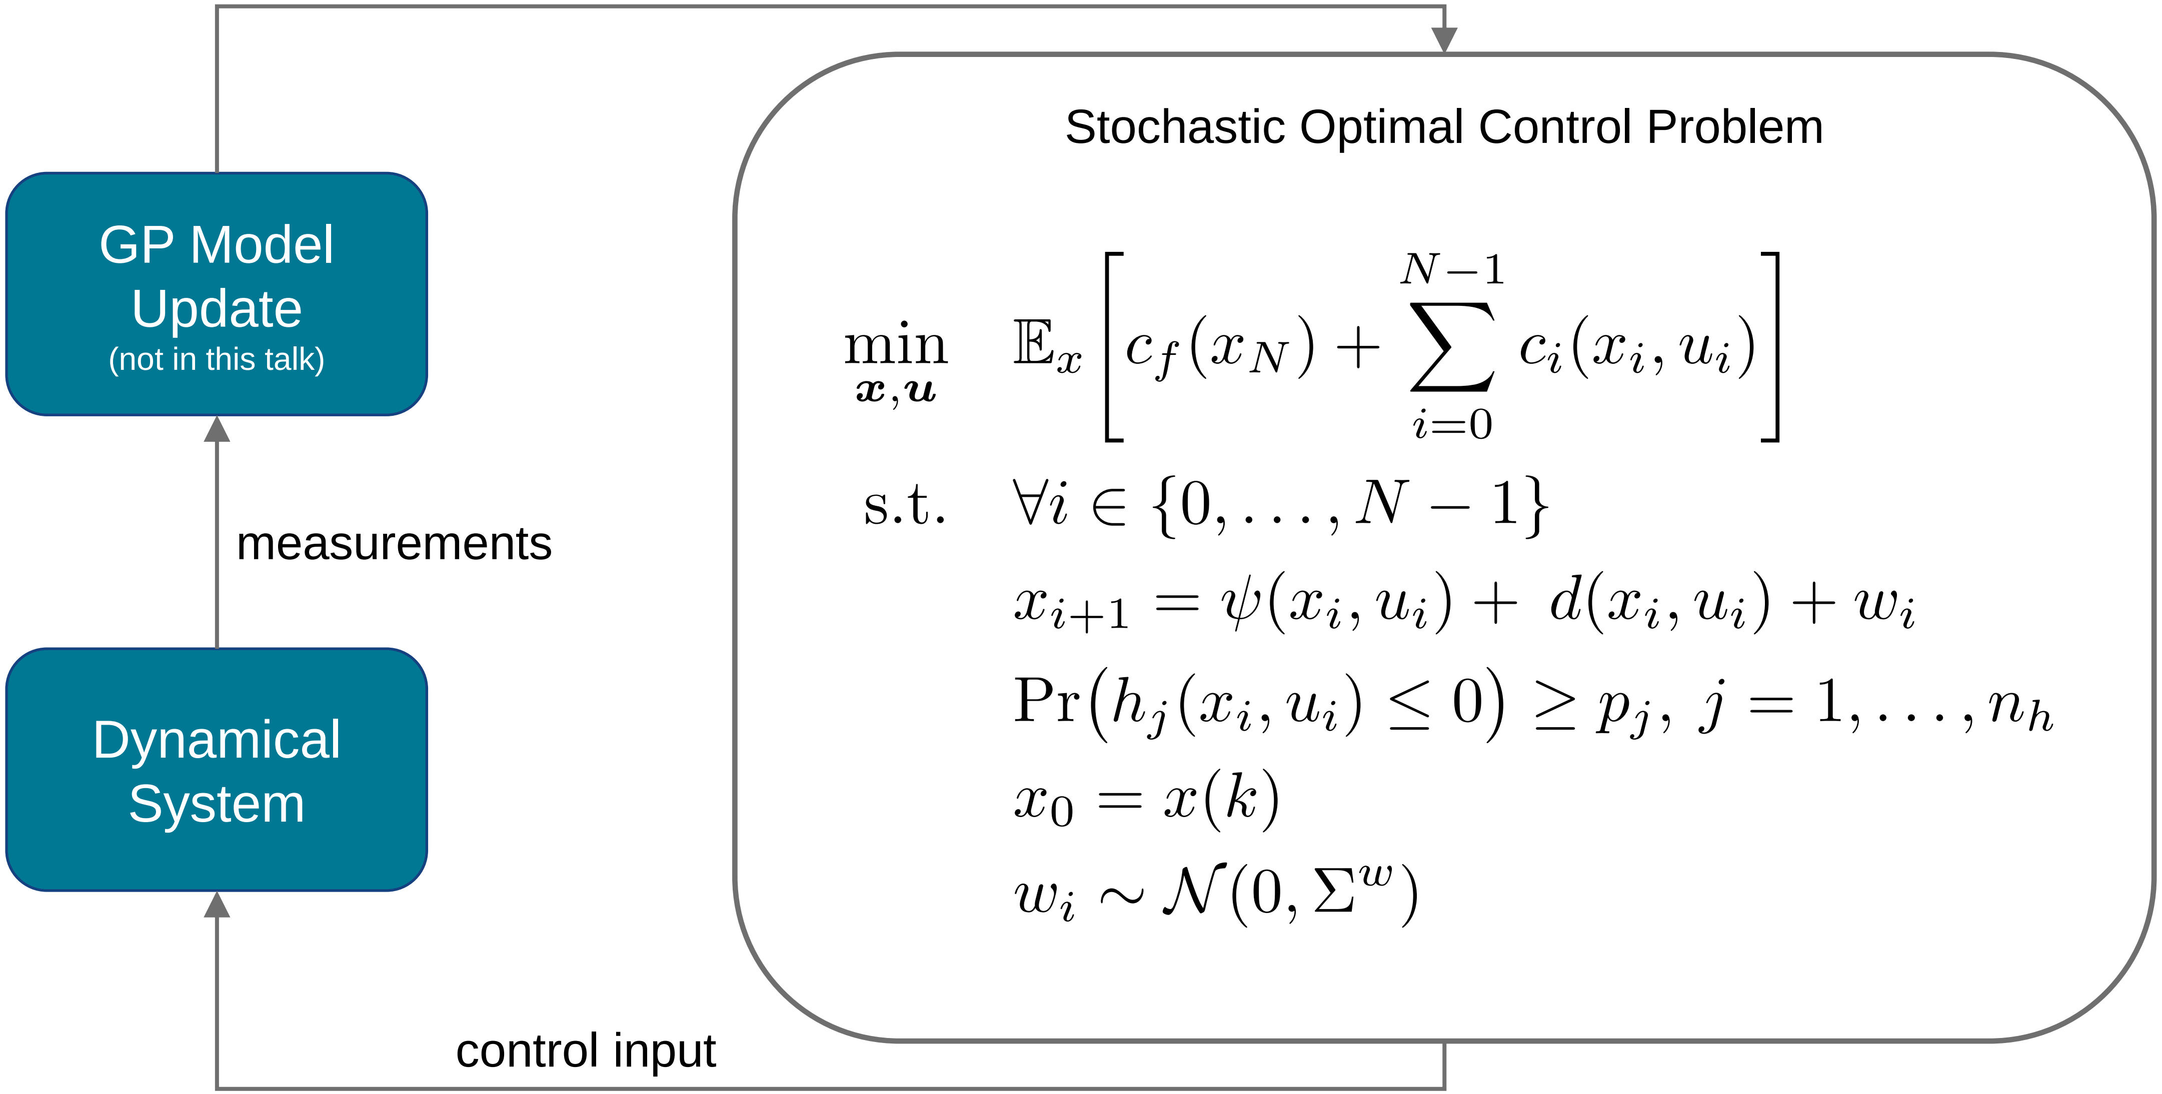

$$
\begin{aligned}
\min_{\boldsymbol{x},\boldsymbol{u}} \quad & \mathbb{E}_x \left[c_f(x_N) + \sum_{i=0}^{N-1} c_i(x_i, u_i) \right] \\
\text{s.t.} \quad & \forall i \in \{0, \dots, N-1\} \\
& x_{i+1} = \psi(x_i, u_i) + d(x_i, u_i) + w_i \\
& \Pr\big(h_j(x_i, u_i) \leq 0\big) \geq p_j,\ j = 1, \dots, n_h \\
& x_0 = x(k) \\
& w_i \sim \mathcal{N}(0, \Sigma^w)
\end{aligned}
$$

GP-MPC – Mean and covariance propagation

Stochastic dynamics model

$$x_{i+1} = \psi(x_i, u_i) + Bd(x_i, u_i) + w_i$$

$$d(x_i, u_i) \sim \mathcal{GP}(\mu^d, \Sigma^d)$$

$$w_i \sim \mathcal{N}(0, \Sigma^w)$$

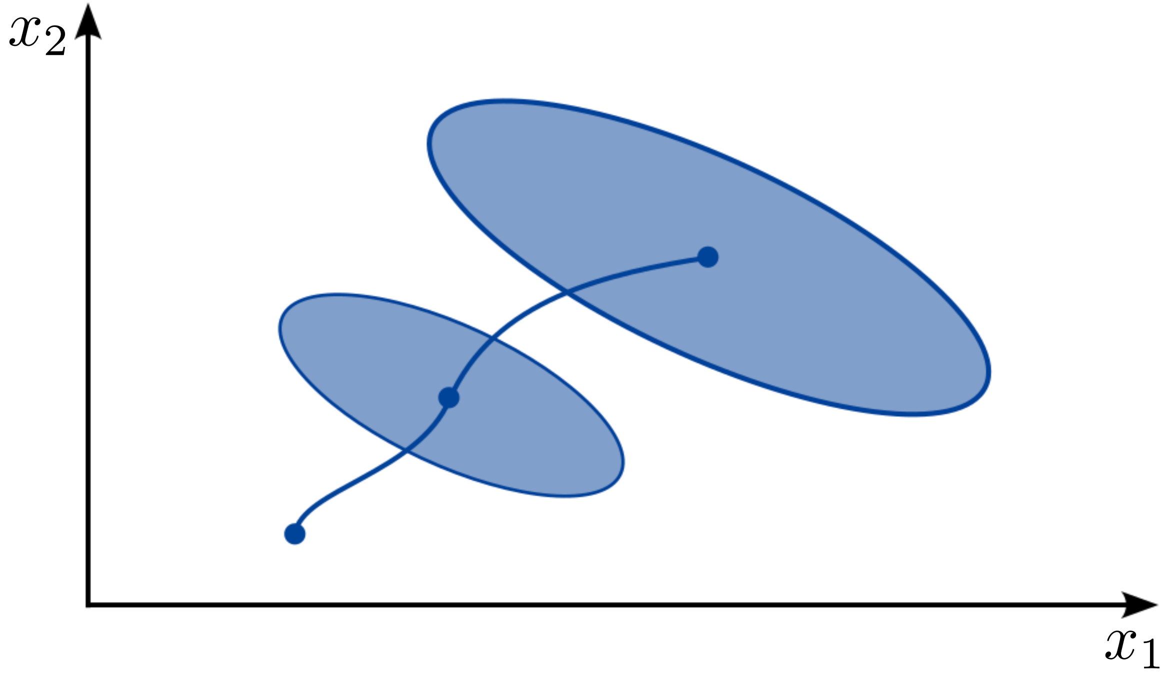

Linearization-based mean and covariance propagation (cf. Girard, Rasmussen, Murray-Smith, 2003)

$$\mu^x_{i+1} = \psi(\mu^x_i, u_i) + B\mu^d(\mu^x_i, u_i)$$

$$\Sigma^x_{i+1} = \tilde{A}_i \Sigma^x_i \tilde{A}_i^\top + B\Sigma^d(\mu^x_i, u_i)B^\top + \Sigma^w \qquad \tilde{A}_i := \frac{\partial}{\partial x}\left(\psi(x, u_i) + B\mu^d(x, u_i)\right)\Big|_{x=\mu^x_i}$$

→ Tube formulation with linear feedback also possible (cf. Hewing et al., 2019, Feng et al., 2020, Messerer et al., 2021)

GP-MPC – Constraint tightening formulation

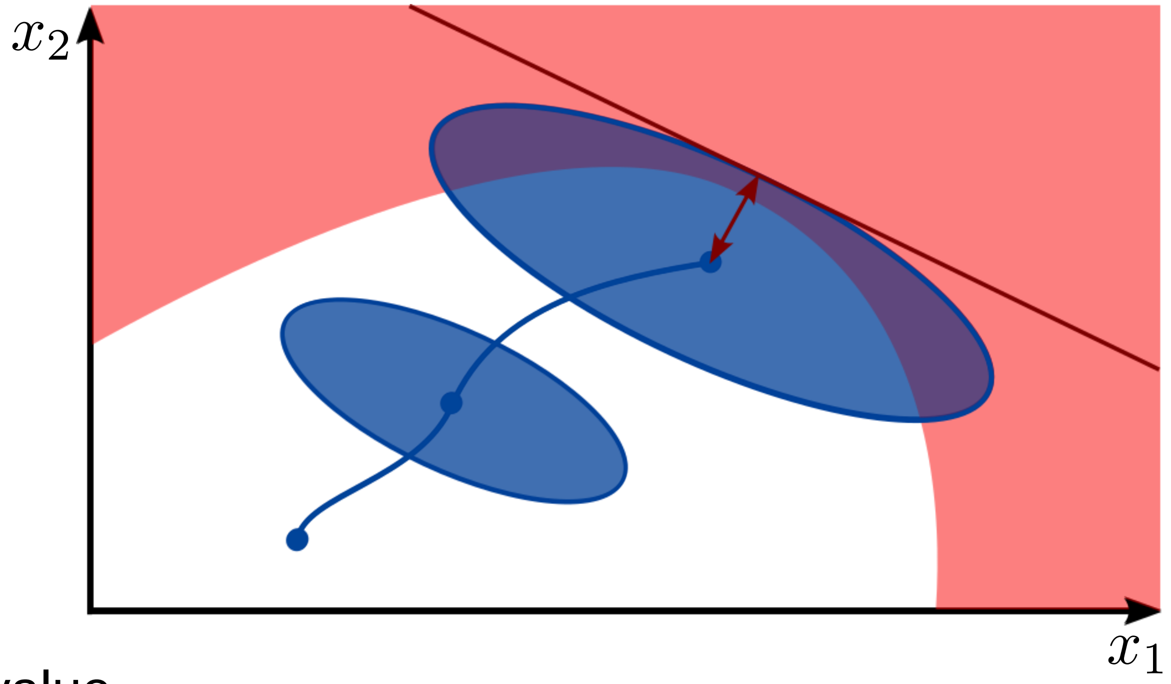

Original chance constraint

$$\Pr(h_j(x_i, u_i) \leq 0) \geq p_j$$

Approximate state distribution as Gaussian

$$x_i \sim \mathcal{N}(\mu_i^x, \Sigma_i^x)$$

Linearization-based tightening around mean value (cf. Srinivasan et. al., 2003)

$$\underbrace{h_j(\mu_i^x, u_i) + \Phi^{-1}(p_j)\sqrt{C_j(\mu_i^x, u_i)\Sigma_i^x C_j(\mu_i^x, u_i)^\top}}_{=: \bar{h}_j(\mu_i^x, u_i, \Sigma_i^x)} \leq 0$$

$$C_j(\mu_i^x, u_i) := \frac{\partial h_j}{\partial x}(\mu_i^x, u_i)$$

GP-MPC – Optimization problem

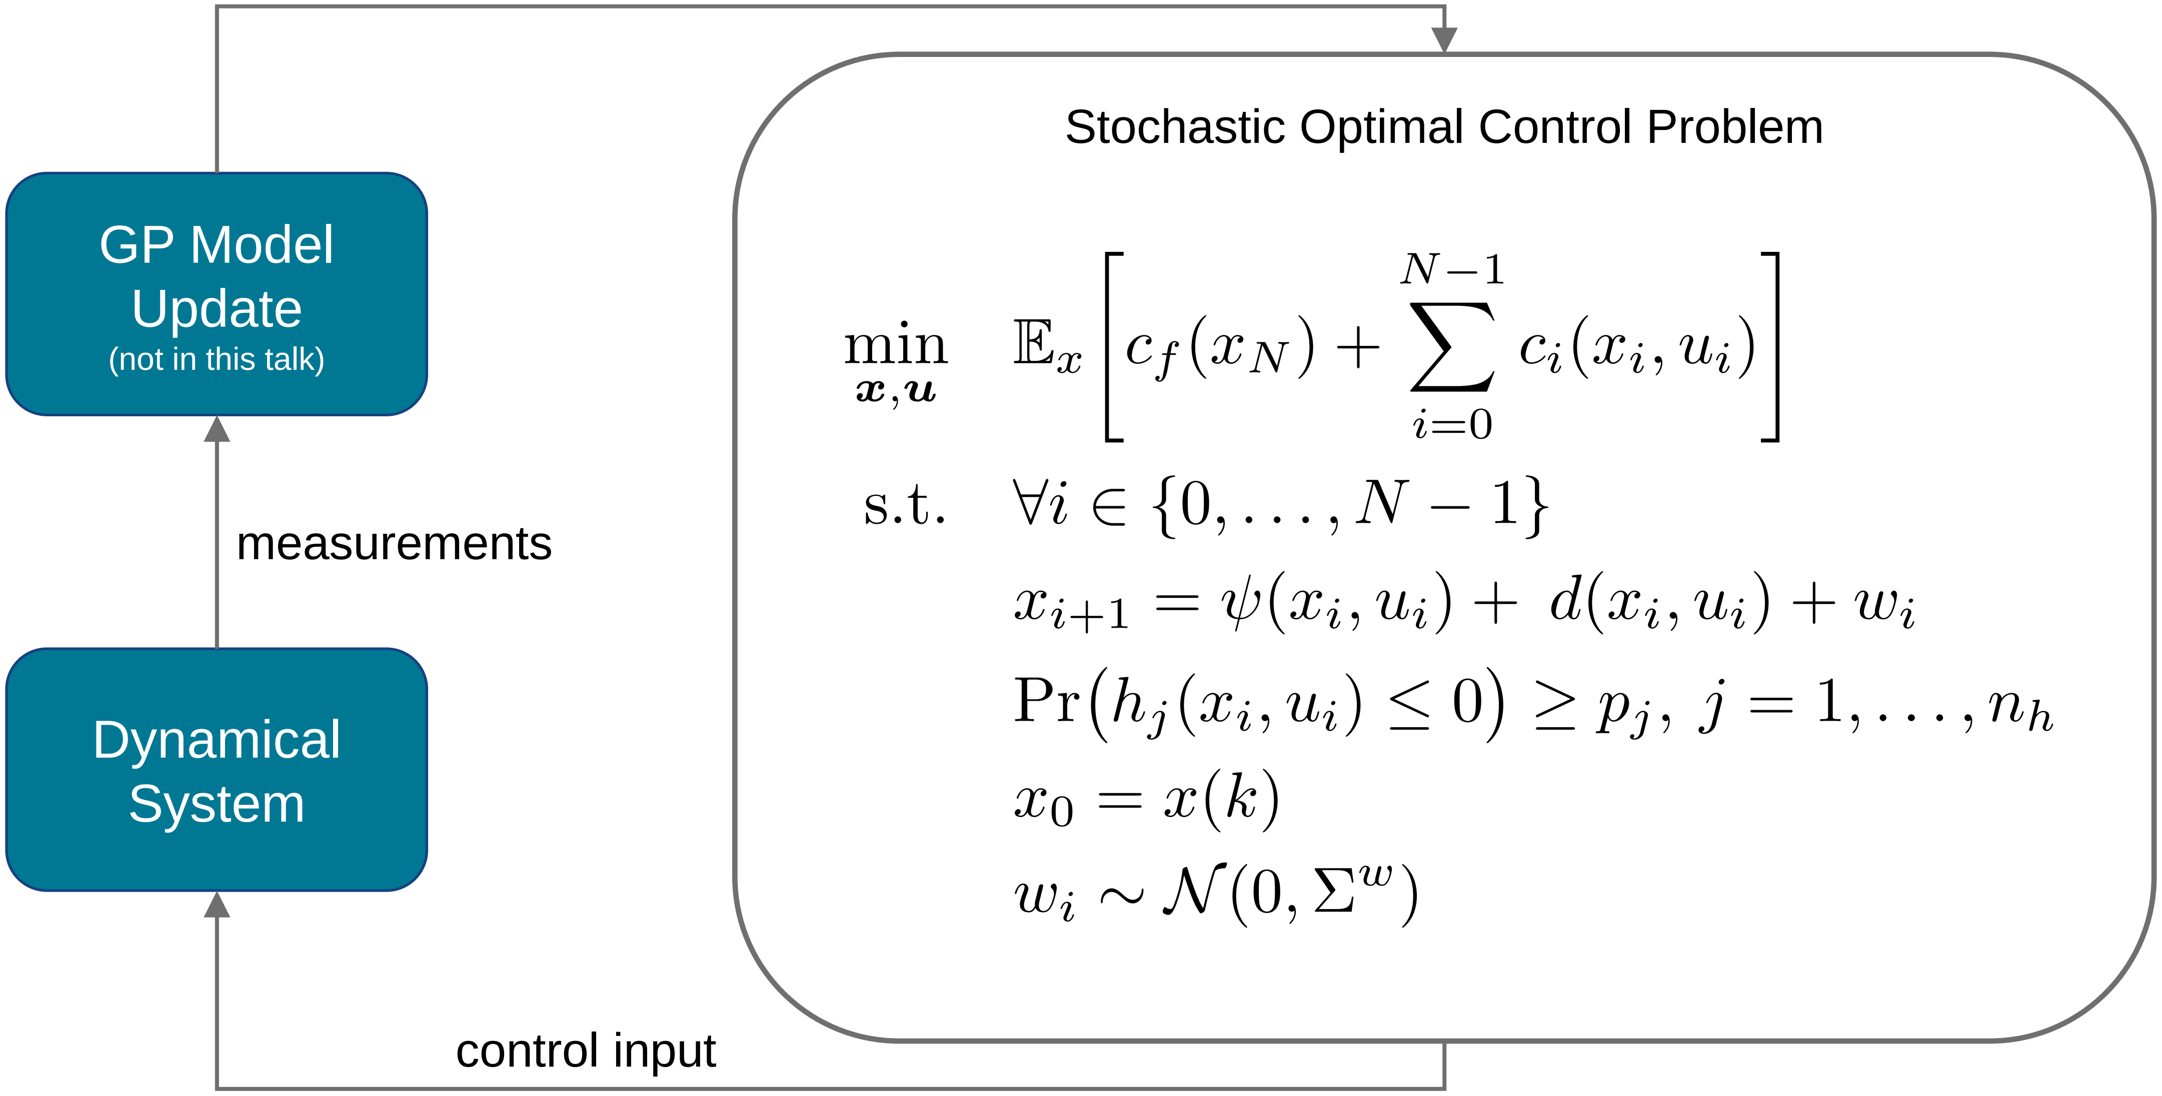

$$
\begin{aligned}
\min_{\boldsymbol{x},\boldsymbol{u}} \quad & \mathbb{E}_x\left[c_f(x_N) + \sum_{i=0}^{N-1} c_i(x_i, u_i)\right] \\
\text{s.t.} \quad & \forall i \in \{0, \ldots, N-1\} \\
& x_{i+1} = \psi(x_i, u_i) + \ d(x_i, u_i) + w_i \\
& \Pr\big(h_j(x_i, u_i) \leq 0\big) \geq p_j, \ j = 1, \ldots, n_h \\
& x_0 = x(k) \\
& w_i \sim \mathcal{N}(0, \Sigma^w)
\end{aligned}
$$

GP-MPC – Optimization Problem

GP Model Update
(not in this talk)

measurements

Dynamical System

control input

GP-MPC Optimal Control Problem [cf. Hewing et al., 2020]

$$
\begin{aligned}
\min_{\boldsymbol{\mu},\boldsymbol{u},\boldsymbol{\Sigma}} \quad & c_f(\mu_N^x) + \sum_{i=0}^{N-1} c_i(\mu_i^x, u_i) \\
\text{s.t.} \quad & \forall i \in \{0, \ldots, N-1\} \\
& \mu_{i+1}^x = \psi(\mu_i^x, u_i) + B\mu^d(\mu_i^x, u_i) \\
& \Sigma_{i+1}^x = \tilde{A}_i \Sigma_i^x \tilde{A}_i^\top + B\Sigma^d(\mu_i^x, u_i)B^\top + \Sigma^w \\
& \bar{h}_j(\mu_i^x, u_i, \Sigma_i^x) \leq 0, \; j = 1, \ldots, n_h \\
& \mu_0^x = x(k) \\
& \Sigma_0^x = 0
\end{aligned}
$$

Expected value approximation

GP-MPC – Computational challenges

Challenges

- Block-sparse structure with $\mathcal{O}(n_x^2)$ **stage variables**
- Gradients of covariance propagation constraints require
  - **Jacobian of GP covariance**
  - **Hessian of GP mean**

→ Complexity $\mathcal{O}(Nn_x^6)$ with sparsity-exploiting interior point solvers

→ Expensive GP evaluations scaling with the number of data points

GP-MPC Optimal Control Problem [cf. Hewing et al., 2020]

$$
\begin{aligned}
\min_{\boldsymbol{\mu},\boldsymbol{u},\boldsymbol{\Sigma}} \quad & c_f(\mu_N^x) + \sum_{i=0}^{N-1} c_i(\mu_i^x, u_i) \\
\text{s.t.} \quad & \forall i \in \{0, \ldots, N-1\} \\
& \mu_{i+1}^x = \psi(\mu_i^x, u_i) + B\mu^d(\mu_i^x, u_i) \\
& \Sigma_{i+1}^x = \tilde{A}_i \Sigma_i^x \tilde{A}_i^\top + B\Sigma^d(\mu_i^x, u_i)B^\top + \Sigma^w \\
& \bar{h}_j(\mu_i^x, u_i, \Sigma_i^x) \leq 0,\ j = 1, \ldots, n_h \\
& \mu_0^x = x(k) \\
& \Sigma_0^x = 0
\end{aligned}
$$

GP-MPC – Computational challenges

Challenges

- Block-sparse structure with $\mathcal{O}(n_x^2)$ **stage variables**
- Gradients of covariance propagation constraints require
  - **Jacobian of GP covariance**
  - **Hessian of GP mean**

$\rightarrow$ Complexity $\mathcal{O}(N n_x^6)$ with sparsity-exploiting interior point solvers

$\rightarrow$ Expensive GP evaluations scaling with the number of data points

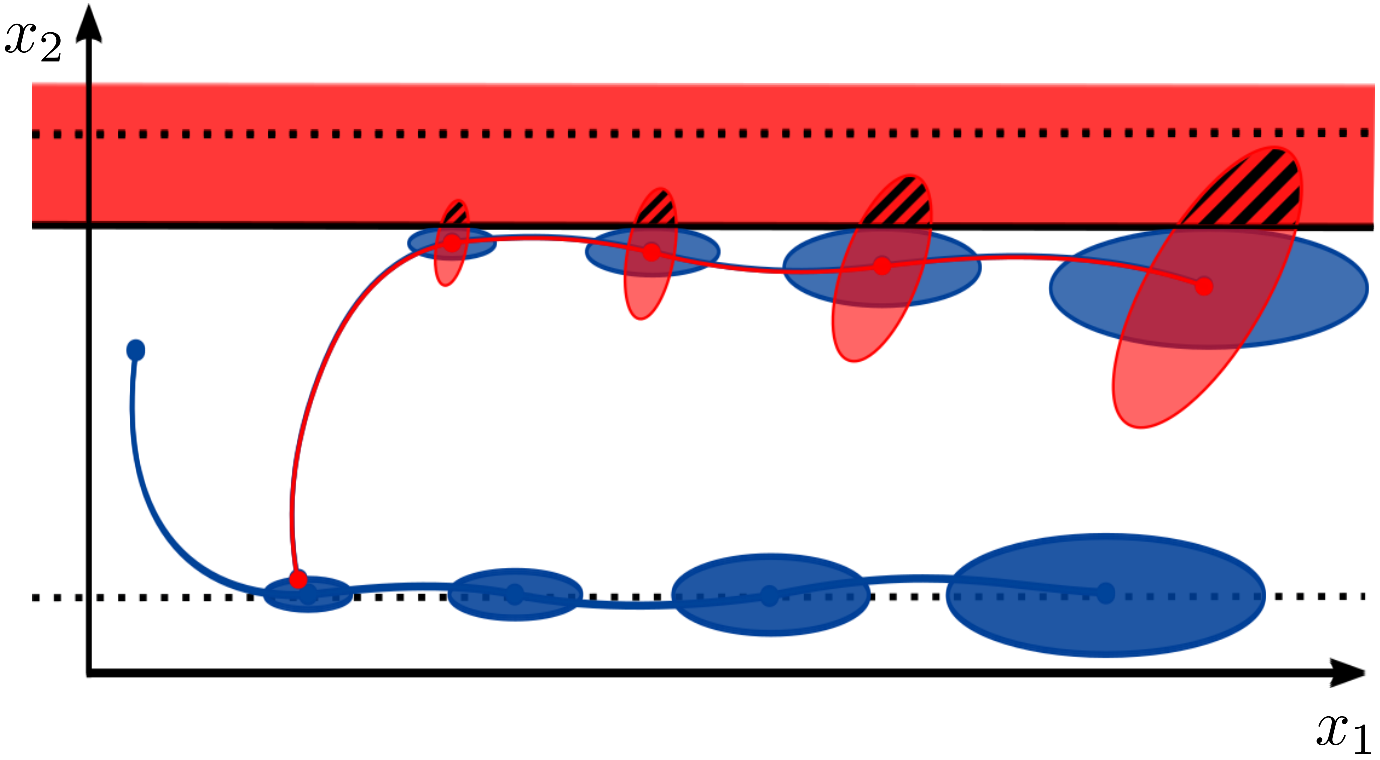

$\rightarrow$ ***Heuristic:*** *Propagate covariances outside the optimizer, based on last MPC solution*
$\rightarrow$ *potentially infeasible solutions!*

GP-MPC – Simplify notation

Simplify notation

y: mean states, inputs

$y := (\mu_0^x, u_0, \ldots, u_{N-1}, \mu_N^x)$

$$
\begin{aligned}
\min_{y,z} \quad & c(y) \\
\text{s.t.} \quad & f(y) = 0 \\
& g(y, z) = 0 \\
& \bar{h}(y, z) \leq 0
\end{aligned}
$$

z: vectorized covariances

GP-MPC Optimal Control Problem [cf. Hewing et al., 2020]

$$
\begin{aligned}
\min_{\boldsymbol{\mu}, \boldsymbol{u}, \boldsymbol{\Sigma}} \quad & c_f(\mu_N^x) + \sum_{i=0}^{N-1} c_i(\mu_i^x, u_i) \\
\text{s.t.} \quad & \forall i \in \{0, \ldots, N-1\} \\
& \mu_{i+1}^x = \psi(\mu_i^x, u_i) + B\mu^d(\mu_i^x, u_i) \\
& \Sigma_{i+1}^x = \tilde{A}_i \Sigma_i^x \tilde{A}_i^\top + B\Sigma^d(\mu_i^x, u_i)B^\top + \Sigma^w \\
& \bar{h}_j(\mu_i^x, u_i, \Sigma_i^x) \leq 0, \; j = 1, \ldots, n_h \\
& \mu_0^x = x(k) \\
& \Sigma_0^x = 0
\end{aligned}
$$

Zero-Order Gaussian Process-based MPC

GP-MPC problem

$$
\begin{aligned}
\min_{y,z} \quad & c(y) \\
\text{s.t.} \quad & f(y) = 0 \\
& g(y, z) = 0 \\
& \bar{h}(y, z) \leq 0
\end{aligned}
$$

Sequential Quadratic Programming

$$
\begin{aligned}
\min_{\Delta y, \Delta z} \quad & \frac{1}{2} \Delta y^\top M_{yy} \Delta y + c_y(\hat{y}) \Delta y \\
\text{s.t.} \quad & 0 = f(\hat{y}) + f_y(\hat{y}) \Delta y \\
& 0 = g(\hat{y}, \hat{z}) + \cancel{g_y(\hat{y}, \hat{z})} \Delta y + g_z(\hat{y}, \hat{z}) \Delta z \\
& 0 \geq \bar{h}(\hat{y}, \hat{z}) + \bar{h}_y(\hat{y}, \hat{z}) \Delta y + \bar{h}_z(\hat{y}, \hat{z}) \Delta z
\end{aligned}
$$

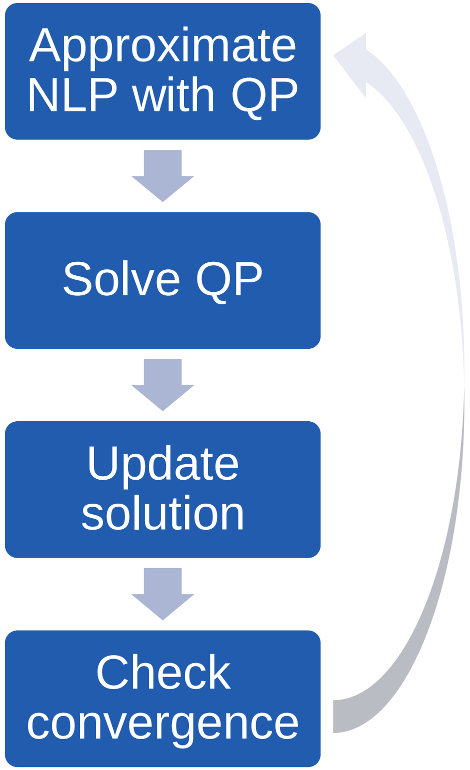

Zero-order approximation *[Bock et al. 2007, Feng et. al. 2020, Zanelli et. al. 2021]*

- Avoids Hessian computation of GP mean and Jacobian computation of GP covariance
- Decouples covariance propagation constraints from QP

Zero-Order Gaussian Process-based MPC

GP-MPC problem

$$\begin{aligned} \min_{y,z} \quad & c(y) \\ \text{s.t.} \quad & f(y) = 0 \\ & g(y,z) = 0 \\ & \bar{h}(y,z) \leq 0 \end{aligned}$$

Sequential Quadratic Programming

$$\Delta z = -g_z(\hat{y},\hat{z})^{-1} g(\hat{y},\hat{z})$$

$$\begin{aligned} \min_{\Delta y} \quad & \frac{1}{2}\Delta y^\top M_{yy} \Delta y + c_y(\hat{y})\Delta y \\ \text{s.t.} \quad & 0 = f(\hat{y}) + f_y(\hat{y})\Delta y \\ & 0 \geq \bar{h}(\hat{y},\hat{z}) + \bar{h}_y(\hat{y},\hat{z})\Delta y + \bar{h}_z(\hat{y},\hat{z})\Delta z \end{aligned}$$

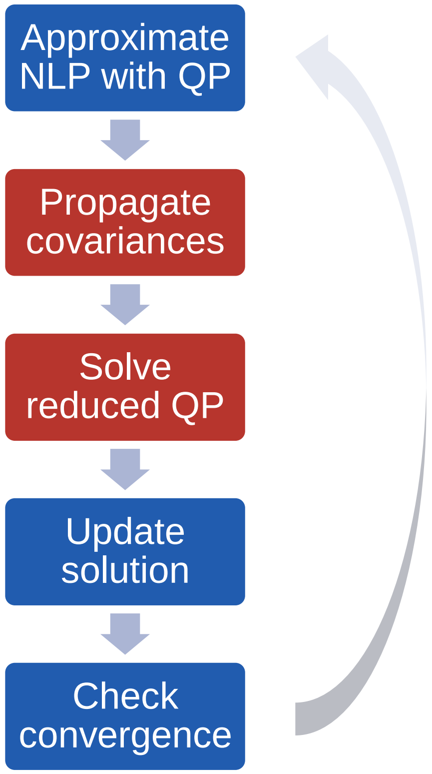

Algorithm properties

- Drastically improved scaling: $\mathcal{O}(n_x^6) \rightarrow \mathcal{O}(n_x^3)$
- Suboptimal, yet feasible solutions
- Preservation of convergence properties for small uncertainties and twice diff. kernels

Zero-Order Gaussian Process-based MPC – Numerical example

Image:
[Kouzoupis et al., 2018]

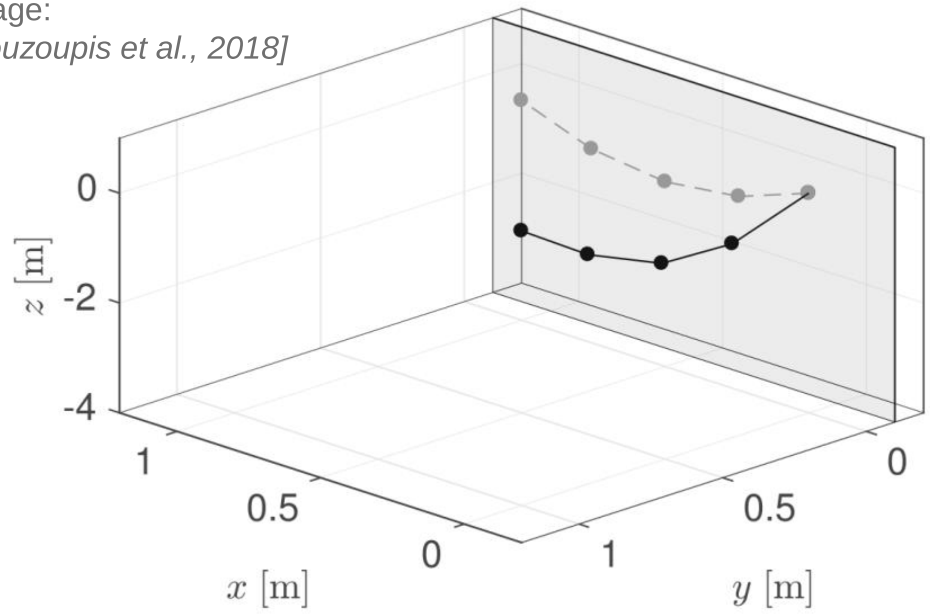

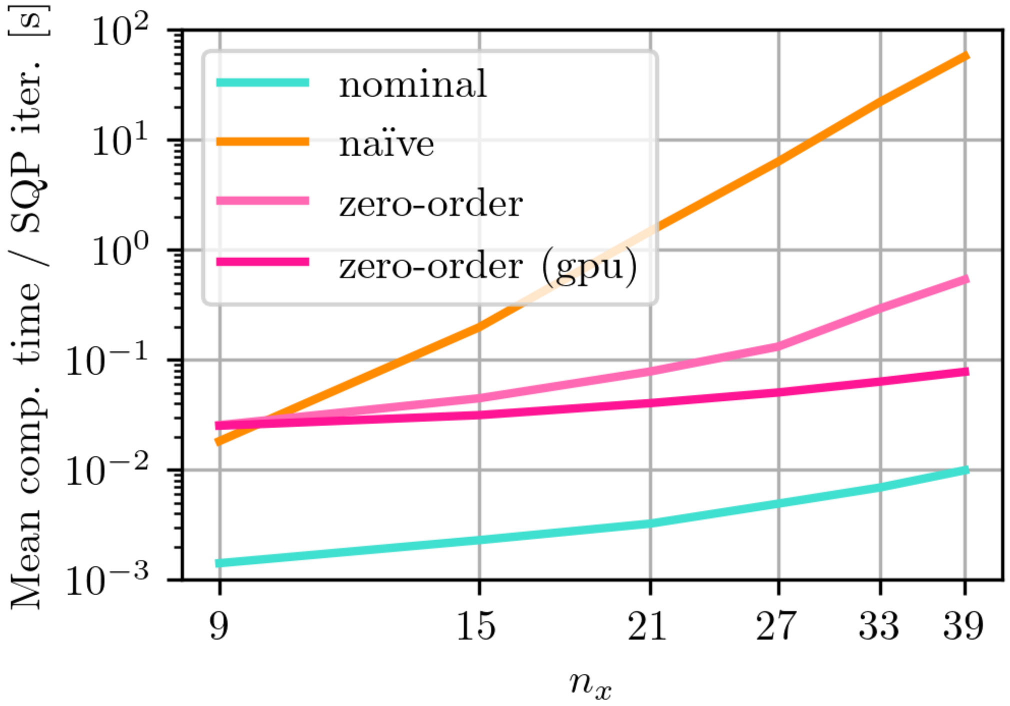

Hanging chain problem *[Wirsching et al., 2006]*

- One end fixed, other end controlled
- Task: Stabilize after strong disturbance, s.t. wall constraint, unknown wind force

- nominal: no uncertainty
- naïve: full optimization with constant uncertainty (no GP data)
- zero-order (gpu): 1500 GP data pts.

Zero-Order Gaussian Process-based MPC – Numerical example

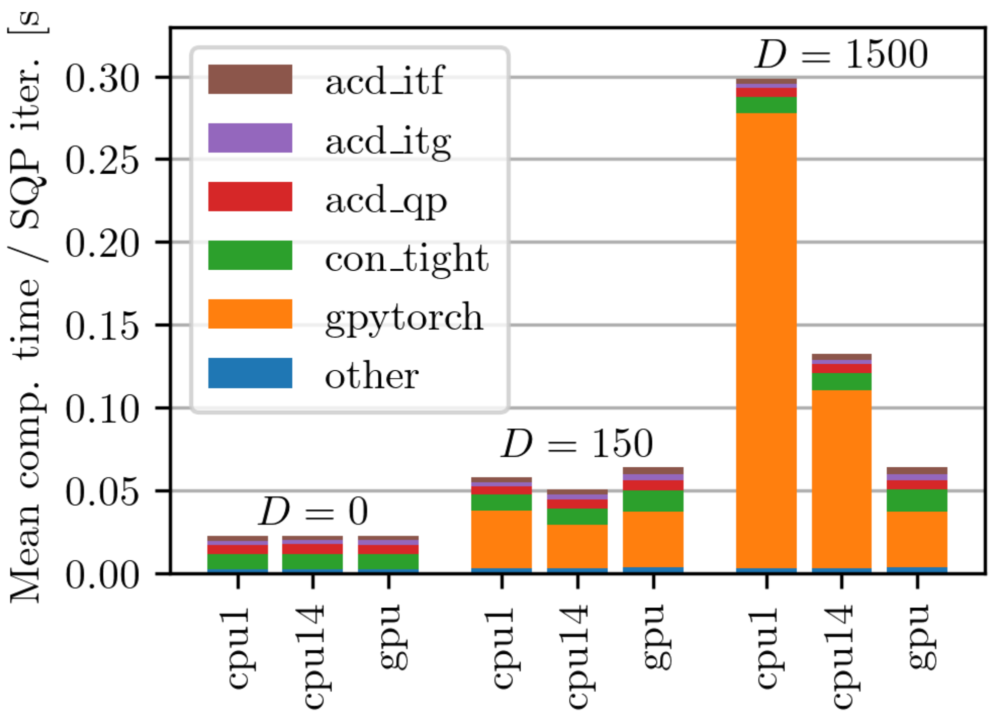

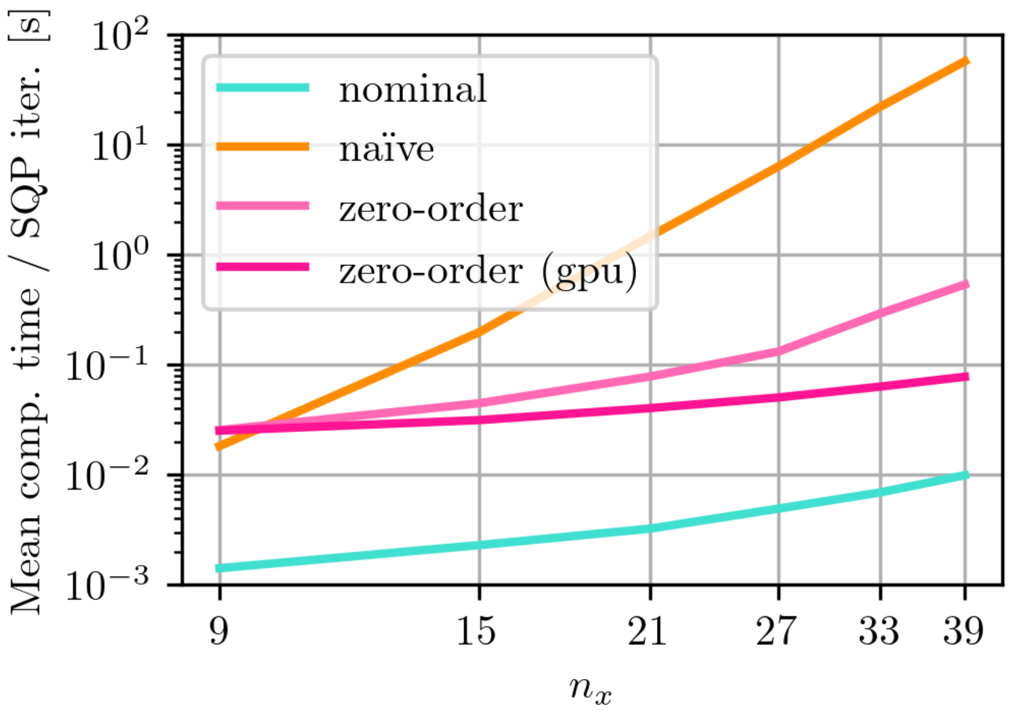

- GP computations = main bottleneck
- GPU parallelization very effective
- Python covar. propagation ~10ms overhead

- nominal: no uncertainty
- naïve: full optimization with constant uncertainty (no GP data)
- zero-order (gpu): 1500 GP data pts.

Summary & Outlook

- Tailored solvers for GP-MPC
  - Significantly improve scaling w.r.t. system dimension
  - Compute **feasible solution at convergence**
  - Preserve SQP algorithm convergence for small uncertainties and smooth kernels
- GPU parallelization + efficient software implementation lead to **~1000x speed-up** for investigated system
- Future work:
  - Combine tailored solver with GP approximation
  - Update GP data and optimize hyper-parameters online
  - Implement algorithm on real hardware

ETH zürich

ELO-X

Funded by the European Union

A. Lahr, A. Zanelli, A. Carron, and M. N. Zeilinger, “Zero-Order Optimization for Gaussian Process-based Model Predictive Control.”, arXiv.2211.15522, 2022.