



Conference Paper

Distributed Model Predictive Safety Certification for Learning-Based Control

Author(s):

Muntwiler, Simon; Wabersich, Kim P.; Carron, Andrea; Zeilinger, Melanie N.

Publication Date:

2020-11

Permanent Link:

<https://doi.org/10.3929/ethz-b-000454995> →

Originally published in:

IFAC-PapersOnLine 53(2), <http://doi.org/10.1016/j.ifacol.2020.12.1205> →

Rights / License:

[Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International](#) →

This page was generated automatically upon download from the [ETH Zurich Research Collection](#). For more information please consult the [Terms of use](#).

Distributed Model Predictive Safety Certification for Learning-based Control^{*}

Simon Muntwiler^{*} Kim P. Wabersich^{*} Andrea Carron^{*}
Melanie N. Zeilinger^{*}

^{*} *Institute for Dynamic Systems and Control, ETH Zürich, Zürich, Switzerland (e-mail: {simonmu,wkim,carrona,mzeilinger}@ethz.ch).*

Abstract: While distributed algorithms provide advantages for the control of complex large-scale systems by requiring a lower local computational load and less local memory, it is a challenging task to design high-performance distributed control policies. Learning-based control algorithms offer promising opportunities to address this challenge, but generally cannot guarantee safety in terms of state and input constraint satisfaction. A recently proposed safety framework for centralized linear systems ensures safety by matching the learning-based input online with the initial input of a model predictive control law capable of driving the system to a terminal set known to be safe. We extend this idea to derive a distributed model predictive safety certification (DMPSC) scheme, which is able to ensure state and input constraint satisfaction when applying any learning-based control algorithm to an uncertain distributed linear system with dynamic couplings. The scheme is based on a distributed tube-based model predictive control (MPC) concept, where subsystems negotiate local tube sizes among neighbors in order to mitigate restrictiveness of the safety approach. In addition, we present a technique for generating a structured ellipsoidal robust positive invariant tube. In numerical simulations, we show that the safety framework ensures constraint satisfaction for an initially unsafe control policy and allows to improve overall control performance compared to robust distributed MPC.

Copyright © 2020 The Authors. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0>)

Keywords: Distributed control, safe learning-based control.

1. INTRODUCTION

One of the key mechanisms in order to improve efficiency in many control applications is the coordination with other systems. Examples include power grids, production processes, or traffic systems. Centralized control of such large-scale complex systems is often computationally infeasible or undesired from a communication point of view (Langbort et al., 2004). Distributed control algorithms address this issue by exploiting the distributed structure of the system using local control inputs that only directly depend on information from neighboring subsystems. In order to achieve optimality of a control algorithm with respect to a given cost metric under such information constraints, learning-based control offers a promising approach. However, learning-based controllers generally fail to guarantee safety in terms of state and input constraint satisfaction, because optimizing the control objective can lead to critical system states or damage to the environment (Amodei et al., 2016), which limits their application to safety-critical real-world systems. Considering the example of power grid control, violations of voltage or current constraints can lead to a reduced lifetime of the grid and are commonly restricted by legislation (Bhattarai et al., 2017).

Recent research has shown advances in the area of learning-based control of centralized systems, as for example teaching a helicopter model to fly aerobatic maneuvers in Abbeel et al. (2010) or a quadrupedal robot to perform advanced locomotion tasks in Hwangbo et al. (2019). The application of centralized learning-based control to multi-agent systems suffers from the curse of dimensionality, since the complexity of the problem scales exponentially with the number of agents (Buşoniu et al., 2010). In this context, distributed learning-based controllers have been investigated, as for example in Sartoretti et al. (2019) for the application to a collective construction task. However, such advances can only be transferred to the control of safety-critical systems once constraint satisfaction can be guaranteed. Safe learning in centralized control systems was intensively studied, e.g., in Fisac et al. (2019), Berkenkamp et al. (2017), but generally does not fit into the information structure of distributed control.

This paper introduces a distributed model predictive safety certification (DMPSC) mechanism for distributed linear systems with couplings in their dynamics and bounded additive disturbances, where the disturbances are essential in order to compensate for uncertainties in the system description. The proposed framework certifies safety of any distributed learning-based control input online by continuously updating a backup trajectory leading towards a set known to be safe, where the first element of the input trajectory equals the current learning-based control input. The trajectory computation is based on

^{*} This work was supported by the Swiss National Science Foundation under grant no. PP00P2_157601/1. The research of Andrea Carron was supported by the Swiss National Centre of Competence in Research NCCR Digital Fabrication (Agreement #51NF40_141853).

distributed model predictive control (MPC) concepts. If no backup trajectory exists, the learning-based input needs to be altered in order to keep the system safe. The MPC safety approach allows for unifying the computation of the backup trajectory and the required modification (Wabersich and Zeilinger, 2018b).

Contributions: The proposed safety framework combines the idea of safety certification through an MPC backup trajectory introduced in Wabersich and Zeilinger (2018b) with robust distributed MPC according to Conte et al. (2013). The distributed MPC scheme computes a nominal state trajectory and applies local feedback to keep the system state within a structured ellipsoidal robust positive invariant (RPI) tube around this nominal trajectory. Considering such a robust MPC approach renders the proposed safety framework capable of handling uncertain model descriptions, which is common in the context of learning, where system models are often derived from data. In order to reduce conservatism when verifying a local learning-based input as being safe to apply, we introduce a negotiation procedure allowing to distribute the size of the local ellipsoidal sets among neighboring subsystems. Further, we describe an adapted synthesis procedure of an RPI tube, which allows to simultaneously compute a structured feedback control law and optimize the shape of the tube.

Related Work: Learning-based control algorithms have been specifically adapted to ensure safety by adding safety constraints in the optimization or adapting the exploration process, see e.g. García and Fernández (2015) for a survey. A more general notion of safety is incorporated by safety frameworks as introduced in Seto et al. (1998), consisting of a set of safe system states and a known safe control law able to keep the system within this set. The model predictive safety certification (MPSC) scheme introduced in Wabersich and Zeilinger (2018b), on which we build in this work, belongs to this category. It scales well in the state dimension and, while the concept is related to the frameworks in Akametalu et al. (2014), Wabersich and Zeilinger (2018a), Fisac et al. (2019), it allows to enlarge any given safe set (Wabersich and Zeilinger, 2018b).

A distributed safety framework based on structured ellipsoidal safe sets was introduced in Larsen et al. (2017). While the approach is less computationally demanding than the proposed method, the ellipsoidal structure of the safe sets can render the certification more conservative, depending on the shape of the state constraints. The distributed safety framework proposed in this work is especially beneficial if the optimal operation point of a system is near to the border of the state constraints.

Structure: We start by introducing the problem and the idea of the MPSC framework in Section 2. In Section 3, we derive our distributed safety framework. Section 4 shows the modified computation method for a distributed ellipsoidal RPI tube. Section 5 gives the results of our numerical simulations and 6 concludes the paper.

2. PROBLEM STATEMENT

2.1 Notation

The set of integers in the interval $[a, \infty) \subseteq \mathbb{N}$ is denoted as $\mathcal{I}_{\geq a}$. The cardinality of a set $\mathcal{I} \subseteq \mathbb{N}$ is denoted as $|\mathcal{I}|$. A stacked vector $v \in \mathbb{R}^n$ consisting of subvectors $v_i \in \mathbb{R}^{n_i}$ with $i \in \mathcal{I} \subseteq \mathbb{N}$ is denoted as $v = \text{col}_{i \in \mathcal{I}}(v_i)$, a block-diagonal matrix $M \in \mathbb{R}^{n \times n}$ consisting of blocks $M_i \in \mathbb{R}^{n_i \times n_i}$ with $i \in \mathcal{I} \subseteq \mathbb{N}$ as $M = \text{diag}_{i \in \mathcal{I}}(M_i)$. The Minkowski sum of two sets $\mathcal{A}_1, \mathcal{A}_2 \subseteq \mathbb{R}^n$ is denoted as $\mathcal{A}_1 \oplus \mathcal{A}_2 = \{a_1 + a_2 \in \mathbb{R}^n | a_1 \in \mathcal{A}_1, a_2 \in \mathcal{A}_2\}$ and the Pontryagin difference as $\mathcal{A}_1 \ominus \mathcal{A}_2 = \{x \in \mathbb{R}^n | x + a_2 \in \mathcal{A}_1, \forall a_2 \in \mathcal{A}_2\}$. Given a set $\mathcal{A} \subseteq \mathbb{R}^n$ and matrix $M \in \mathbb{R}^{l \times n}$ we define $M\mathcal{A} = \{x \in \mathbb{R}^l | x = Ma, a \in \mathcal{A}\}$.

2.2 Distributed Linear Systems

We consider a network of $M \in \mathbb{N}$ time-invariant coupled linear subsystems with discrete-time dynamics

$$x_i(t+1) = \left(\sum_{j=1}^M A_{ij} x_j(t) \right) + B_i u_i(t) + G_i w_i(t), \quad (1)$$

where $x_i(t) \in \mathbb{R}^{n_i}$, $u_i(t) \in \mathbb{R}^{m_i}$ and $w_i(t) \in \mathbb{R}^{p_i}$ are the state, input and bounded disturbance of subsystem i at time step t , respectively, $A_{ij} \in \mathbb{R}^{n_i \times n_j}$, $B_i \in \mathbb{R}^{n_i \times m_i}$, and $G_i \in \mathbb{R}^{n_i \times p_i}$. We denote the set of indices of all subsystems as $\mathcal{M} = \{1, \dots, M\}$. The set of neighbors \mathcal{N}_i of subsystem i contains index i itself, as well as all indices of subsystems j , for which A_{ij} includes entries not equal to zero. The local system dynamics of subsystem i can be written as

$$x_i(t+1) = A_{\mathcal{N}_i} x_{\mathcal{N}_i}(t) + B_i u_i(t) + G_i w_i(t), \quad (2)$$

where $A_{\mathcal{N}_i} \in \mathbb{R}^{n_i \times n_{\mathcal{N}_i}}$ and $x_{\mathcal{N}_i}(t) = \text{col}_{j \in \mathcal{N}_i}(x_j(t)) \in \mathbb{R}^{n_{\mathcal{N}_i}}$. Each subsystem i is subject to polytopic state and input constraints

$$\mathcal{X}_{\mathcal{N}_i} = \{x_{\mathcal{N}_i} \in \mathbb{R}^{n_{\mathcal{N}_i}} | H_{\mathcal{N}_i} x_{\mathcal{N}_i} \leq h_{\mathcal{N}_i}\} \subseteq \mathbb{R}^{n_{\mathcal{N}_i}}, \quad (3a)$$

$$\mathcal{U}_i = \{u_i \in \mathbb{R}^{m_i} | O_i u_i \leq o_i\} \subseteq \mathbb{R}^{m_i}, \quad (3b)$$

where $H_{\mathcal{N}_i} \in \mathbb{R}^{q_{\mathcal{N}_i} \times n_{\mathcal{N}_i}}$, $h_{\mathcal{N}_i} \in \mathbb{R}^{q_{\mathcal{N}_i}}$, $O_i \in \mathbb{R}^{r_i \times m_i}$ and $o_i \in \mathbb{R}^{r_i}$. Both state and input constraints contain the origin in their interior. The local disturbances $w_i(t)$ are contained in the ellipsoidal sets $\mathcal{W}_i = \{w_i \in \mathbb{R}^{p_i} | w_i^\top Q_i w_i \leq q_i\}$ with $Q_i \in \mathbb{R}^{p_i \times p_i}$ and $q_i \in \mathbb{R}$. The global dynamics of the network is described by

$$x(t+1) = Ax(t) + Bu(t) + Gw(t), \quad (4)$$

where $A \in \mathbb{R}^{n \times n}$, $B = \text{diag}_{i \in \mathcal{M}}(B_i) \in \mathbb{R}^{n \times m}$, $G = \text{diag}_{i \in \mathcal{M}}(G_i) \in \mathbb{R}^{n \times p}$, $x(t) = \text{col}_{i \in \mathcal{M}}(x_i(t)) \in \mathbb{R}^n$ is the global state, $u(t) = \text{col}_{i \in \mathcal{M}}(u_i(t)) \in \mathbb{R}^m$ the global input and $w(t) = \text{col}_{i \in \mathcal{M}}(w_i(t)) \in \mathbb{R}^p$ the global bounded additive disturbance. The global state, input and disturbance are subject to the convex compact constraint sets

$$\mathcal{X} = \{x \in \mathbb{R}^n | Hx \leq h\}, \quad (5a)$$

$$\mathcal{U} = \mathcal{U}_1 \times \dots \times \mathcal{U}_M = \{u \in \mathbb{R}^m | Ou \leq o\}, \quad (5b)$$

$$\mathcal{W} = \mathcal{W}_1 \times \dots \times \mathcal{W}_M, \quad (5c)$$

where $H \in \mathbb{R}^{q \times n}$, $h \in \mathbb{R}^q$, $O = \text{diag}_{i \in \mathcal{M}}(O_i) \in \mathbb{R}^{r \times m}$ and $o = \text{col}_{i \in \mathcal{M}}(o_i) \in \mathbb{R}^r$. We assume the system matrices (A, B) to be stabilizable and the system state to be fully measurable.

Remark 2.1. Considering system models with additive disturbances allows to handle systems with external disturbances and systems with uncertainty in the system parameters. It also allows to simplify even perfectly known

system models, e.g., by reducing the size of the state space or neglecting nonlinearities, in order to reduce the computational load of solving the problem in a distributed fashion. Furthermore, it would allow for inexact primal dual optimization according to Köhler et al. (2019).

2.3 Communication and Distributed Optimization

Bidirectional communication among subsystems i and j is possible if $i \in \mathcal{N}_j$ or $j \in \mathcal{N}_i$. We further assume the underlying communication graph to be connected, meaning there is a communication path between any two subsystems. In order to apply the DMPSC scheme in a distributed manner, we require a distributed optimization algorithm which only uses direct communication between neighboring subsystems. An overview of parallel and distributed computation is given in Bertsekas and Tsitsiklis (1989).

2.4 Model Predictive Safety Certification

In the following, we briefly review the ideas behind the MPSC scheme introduced in Wabersich and Zeilinger (2018b), which serves as the basis for the presented distributed safety framework. The MPSC scheme guarantees safety in terms of state and input constraint satisfaction for a linear system (4) with constraints (5). The idea is to find a safe set of states $\mathcal{S} \subseteq \mathcal{X}$, for which we know a safe backup control law $u_{\mathcal{B}}$ for every state in the set. This backup control law is able to keep the system state within \mathcal{S} and therefore ensures constraint satisfaction. A learning-based input $u_{\mathcal{L}}$ is applied to the system if it keeps the system state within \mathcal{S} . Otherwise, we can use the safe backup control law $u_{\mathcal{B}}$. The safe set \mathcal{S} and safety control law $u_{\mathcal{S}}$ are formally introduced in the following definition.

Definition 2.1. A set $\mathcal{S} \subseteq \mathcal{X}$ is called a safe set for system (4) if there exists a known safe backup control law $u_{\mathcal{B}} : \mathbb{R}^n \times \mathbb{R}^m \times \mathcal{I}_{\geq 0} \rightarrow \mathcal{U}$ such that for an arbitrary policy $u_{\mathcal{L}} : \mathcal{I}_{\geq 0} \rightarrow \mathbb{R}^m$, the application of the safety control law

$$u_{\mathcal{S}}(t) := \begin{cases} u_{\mathcal{L}}(t), & \text{if } u_{\mathcal{L}} \in \mathcal{U} \wedge \{Ax + Bu_{\mathcal{L}}\} \oplus GW \subseteq \mathcal{S} \\ u_{\mathcal{B}}(x(t), u_{\mathcal{L}}(t), t), & \text{otherwise} \end{cases}$$

guarantees that the system state $x(t)$ is contained in $\mathcal{S} \subseteq \mathcal{X}$ for all $t \geq \bar{t}$ if $x(\bar{t}) \in \mathcal{S}$.

In the MPSC scheme, the safe set \mathcal{S} is implicitly described as the set of states where we know an MPC trajectory leading the system to a safe terminal set $\mathcal{S}_f \subseteq \mathbb{R}^n$, where constraint satisfaction can be guaranteed for all times. Robust MPC according to Mayne et al. (2005) is used, in order to account for uncertainty in the system model. The scheme can be formulated as one MPSC optimization problem, which, for a given system state $x(t)$ and learning-based input $u_{\mathcal{L}}(t)$, optimizes over an MPC trajectory leading the system to the terminal set. The objective of the problem is to minimize the ‘distance’ between the learning-based input and the safety control law $u_{\mathcal{S}}(t)$. If possible, $u_{\mathcal{S}}(t)$ is chosen equal to $u_{\mathcal{L}}(t)$, otherwise a minimal safety-ensuring modification is computed.

The goal in this paper is to ensure constraint satisfaction of an uncertain distributed linear system (1) while applying a distributed control law $u_{\mathcal{L}_i}$, for which we extend the MPSC scheme to distributed systems in the following. In this

paper, we assume $u_{\mathcal{L}_i}$ to be learning-based. In principle, $u_{\mathcal{L}_i}$ could represent any control policy, e.g., human control inputs.

3. DISTRIBUTED MODEL PREDICTIVE SAFETY CERTIFICATION

The distribution of the MPSC scheme is based on robust distributed MPC as introduced in Conte et al. (2013). Among the main challenges in order to transfer the centralized MPSC scheme to the distributed case are ensuring recursive feasibility and obtaining an RPI set in order to account for disturbances. In this section, we will first introduce the robust distributed MPC backup control law. Then we state the distributed safety certification problem and introduce the proposed negotiation procedure among subsystems and the form of the terminal constraint. We conclude the section by giving the theoretical safety guarantee of our distributed framework.

3.1 Robust Distributed Backup Control Law

The robust distributed MPC backup control law aims at keeping the distributed system (1) near to the nominal system

$$z_i(k+1|t) = A_{\mathcal{N}_i} z_{\mathcal{N}_i}(k|t) + B_i v_i(k|t) \quad \forall k \in \mathcal{I}_{\geq 0}, \quad (6)$$

where $z_i(k|t) \in \mathbb{R}^{n_i}$, $z_{\mathcal{N}_i}(k|t) = \text{col}_{j \in \mathcal{N}_i}(z_j(k|t)) \in \mathbb{R}^{n_{\mathcal{N}_i}}$ and $v_i(k|t) \in \mathbb{R}^{m_i}$ are the nominal state, neighborhood state and input of subsystem i at prediction step k computed at time step $t \in \mathbb{Z}$. We denote the deviation between real and nominal local system state as $e_i(k|t) = x_i(t+k) - z_i(k|t) \in \mathbb{R}^{n_i}$, and similarly $e_{\mathcal{N}_i}(k|t) \in \mathbb{R}^{n_{\mathcal{N}_i}}$ and $e(k|t) \in \mathbb{R}^n$ for the neighborhood and global deviation.

Assumption 3.1. There exists a stabilizing linear state feedback control law with distributed structure for system (4) of the form $u = K_{\Omega} x = \text{col}_{i \in \mathcal{M}}(K_{\Omega, i} x_{\mathcal{N}_i})$, where $K_{\Omega} \in \mathbb{R}^{m \times n}$ and $K_{\Omega, i} \in \mathbb{R}^{m_i \times n_{\mathcal{N}_i}}$.

Based on K_{Ω} , the tube-based control law

$$\begin{aligned} u(k|t) &= v(k|t) + K_{\Omega} (x(t+k) - z(k|t)) \\ &= v(k|t) + \text{col}_{i \in \mathcal{M}}(K_{\Omega, i} (x_{\mathcal{N}_i}(t+k) - z_{\mathcal{N}_i}(k|t))) \end{aligned} \quad (7)$$

keeps the real system state within a robust positive invariant (RPI) tube around the nominal system state.

Definition 3.1. A set $\Omega \subseteq \mathbb{R}^n$ is called a robust positively invariant (RPI) set for the error dynamics if

$$e(0|t) \in \Omega \Rightarrow e(k|t) \in \Omega \quad (8)$$

for all $k \in \mathcal{I}_{\geq 0}$ and all disturbances $w(k) \in \mathcal{W}$.

In order to obtain a distributed problem structure we consider ellipsoidal RPI sets. An ellipsoidal RPI set of the form $\Omega = \{e \in \mathbb{R}^n | e^{\top} P e \leq 1\}$, with $P \in \mathbb{R}^{n \times n}$ needs to fulfill the implication

$$\left. \begin{aligned} e(k|t)^{\top} P e(k|t) &\leq 1 \\ w_i(t+k)^{\top} Q_i w_i(t+k) &\leq q_i \forall i \in \mathcal{M} \end{aligned} \right\} \Rightarrow e(k+1|t)^{\top} P e(k+1|t) \leq 1, \quad (9)$$

which is equivalent to (19) in Conte et al. (2013). The following defines a structured RPI set.

Definition 3.2. An ellipsoidal RPI set $\Omega \subseteq \mathbb{R}^n$ is structured if $\Omega = \{e \in \mathbb{R}^n | \sum_{i=1}^M e_{\mathcal{N}_i}^{\top} P_{\mathcal{N}_i} e_{\mathcal{N}_i} \leq 1\}$ with all $P_{\mathcal{N}_i} \in \mathbb{R}^{n_{\mathcal{N}_i} \times n_{\mathcal{N}_i}}$ positive semi-definite, i.e., if the shape matrix has a distributed structure.

Given such a structured RPI set Ω , we can define local sets $\Omega_{\mathcal{N}_i} = \{e_{\mathcal{N}_i} \in \mathbb{R}^{n_{\mathcal{N}_i}} | e_{\mathcal{N}_i}^\top P_{\mathcal{N}_i} e_{\mathcal{N}_i} \leq \beta_i\}$. The global condition $e(k|t) \in \Omega$ is equivalent to

$$\forall i \in \mathcal{M} \quad \exists \quad \beta_i \geq 0 : \quad e_{\mathcal{N}_i}(k|t) \in \Omega_{\mathcal{N}_i}, \sum_{i=1}^M \beta_i \leq 1 \quad (10)$$

which follows by summing up the inequalities of the local sets $\Omega_{\mathcal{N}_i}$.

Equation (23) in Conte et al. (2013) states the synthesis problem to obtain a structured ellipsoidal RPI set according to Definition 3.2. The objective of the problem can be chosen to find an RPI set with minimal volume, provided a structured control law according to Assumption 3.1 is known. A possible way to find such a control law is described in Conte et al. (2016). In Section 4 we provide a design method that allows for simultaneously finding a distributed feedback control law according to Assumption 3.1, which can importantly reduce conservatism.

The set Ω is used to obtain the tightened constraint sets $\bar{\mathcal{X}} = \mathcal{X} \ominus \Omega$ and $\bar{\mathcal{U}} = \mathcal{U} \ominus K_\Omega \Omega$ for the nominal system in order to ensure $(x, u) \in \mathcal{X} \times \mathcal{U}$ when applying (7). The structure of Ω can be exploited to perform the constraint tightening by distributed optimization as introduced in Conte et al. (2013), resulting in local tightened constraint sets $\bar{\mathcal{X}}_{\mathcal{N}_i}$ and $\bar{\mathcal{U}}_i$.

3.2 Distributed Safety Certification Problem

Using the robust distributed MPC control law we can write the DMPSC problem

$$\min_{\mathbf{z}, \mathbf{v}, \tilde{\mathbf{u}}, \tilde{\beta}, \Delta\beta} \sum_{i=1}^M \|u_{\mathcal{L}_i} - \tilde{u}_i\| \quad (11a)$$

s.t. $\forall i \in \mathcal{M} :$

$$x_{\mathcal{N}_i}(t) \in z_{\mathcal{N}_i}(0|t) \oplus \Omega_{\mathcal{N}_i}(t), \quad (11b)$$

$$\sum_{j \in \mathcal{N}_i \setminus i} \frac{\Delta\beta_j}{|\mathcal{N}_j| - 1} = 0 \quad (11c)$$

$$\tilde{\beta}_i = \beta_i(t) + \Delta\beta_i \quad (11d)$$

$$\tilde{\beta}_i \geq 0 \quad (11e)$$

$\forall k \in \{0, \dots, N-1\} :$

$$z_i(k+1|t) = A_{\mathcal{N}_i} z_{\mathcal{N}_i}(k|t) + B_i v_i(k|t) \quad (11f)$$

$$(z_{\mathcal{N}_i}(k|t), v_i(k|t)) \in \bar{\mathcal{X}}_{\mathcal{N}_i} \times \bar{\mathcal{U}}_i \quad (11g)$$

$$z_i(N|t) \in \mathcal{X}_{f,i}(\alpha_i(t)) \quad (11h)$$

$$\tilde{u}_i = v_i(0|t) + K_{\Omega,i}(x_{\mathcal{N}_i}(t) - z_{\mathcal{N}_i}(0|t)), \quad (11i)$$

where $\tilde{\mathbf{u}} = \text{col}_{i \in \mathcal{M}}(\tilde{u}_i)$, $\mathbf{z} = \text{col}_{k \in \{0, \dots, N\}}(z(k|t))$ and $\mathbf{v} = \text{col}_{k \in \{0, \dots, N-1\}}(v(k|t))$. The objective of the optimization is to find a distributed control input \tilde{u}_i as ‘close’ as possible to a distributed learning-based input $u_{\mathcal{L}_i}$. Constraint (11b) ensures that the system state lies within a tube around the initial nominal state, (11f) enforces the nominal dynamics and (11g) the tightened nominal constraints, (11h) ensures that the final nominal state lies within the terminal set and finally (11i) defines \tilde{u}_i to be a tube-based control law. We will discuss the negotiation procedure behind constraints (11c)-(11e) and the details of the terminal constraint (11h) in the following subsections. We denote the feasible set of (11) by

$$\mathcal{X}_{\mathcal{N}} := \{x \in \mathbb{R}^n | (11b) - (11i) \quad \forall i \in \mathcal{M}\} \subseteq \mathcal{X}. \quad (12)$$

The resulting certified distributed control law for each subsystem i is given by

$$\kappa_i(x_{\mathcal{N}_i}(t)) = \tilde{u}_i^*(x_{\mathcal{N}_i}(t)), \quad (13)$$

where $*$ denotes the optimal solution of the problem (11).

3.3 Negotiation

In order to ensure recursive feasibility, robust MPC according to Mayne et al. (2005) constrains the initial nominal state such that the state of the system lies in a tube around it, i.e., $x(t) \in z(0|t) \oplus \Omega$. For the distributed problem in Conte et al. (2013), this is enforced by constraint (11b) based on local sets $\Omega_{\mathcal{N}_i}(t) = \{e_{\mathcal{N}_i}(0|t) \in \mathbb{R}^{n_{\mathcal{N}_i}} | e_{\mathcal{N}_i}(0|t)^\top P_{\mathcal{N}_i} e_{\mathcal{N}_i}(0|t) \leq \tilde{\beta}_i(t)\}$ with $\tilde{\beta}_i(t) : \mathbb{N} \rightarrow [0, 1]$ and $\sum_{i=1}^M \tilde{\beta}_i(t) \leq 1$. The main goal for the DMPSC scheme is to choose the distributed control input \tilde{u}_i as ‘close’ as possible to the learning-based input $u_{\mathcal{L}_i}$. Since \tilde{u}_i depends on $z_{\mathcal{N}_i}(0|t)$, we want the constraint (11b) to be as little restrictive as possible. Ideally, we would therefore add the constraint $\sum_{i=1}^M \tilde{\beta}_i(t) \leq 1$ to the optimization problem (11), which is, however, a centralized condition. In order to enable distributed optimization, we introduce $\tilde{\beta}_i(t)$ as

$$\tilde{\beta}_i(t) = \beta_i(t) + \Delta\beta_i \quad (14)$$

where $\beta_i(t)$ is a predefined function satisfying $\sum_{i=1}^M \beta_i(t) \leq 1$ and the value of $\Delta\beta_i$ can be negotiated with neighboring subsystems, while ensuring that $\sum_{i=1}^M \Delta\beta_i = 0$ in a distributed manner. By definition, this guarantees that $\sum_{i=1}^M \tilde{\beta}_i(t) \leq 1$. This formulation enables larger values $\tilde{\beta}_i(t)$ for subsystems, where we need more flexibility to choose $z_{\mathcal{N}_i}(0|t)$ in order to represent a learning-based input.

A candidate solution for $\beta_i(t)$ introduced in Conte et al. (2013) is

$$\beta_i(t) := e_{\mathcal{N}_i}^*(1|t-1)^\top P_{\mathcal{N}_i} e_{\mathcal{N}_i}^*(1|t-1) \quad \forall i \in \mathcal{M} \quad (15)$$

where $e_{\mathcal{N}_i}^*(k|t) = x_{\mathcal{N}_i}(t+k) - z_{\mathcal{N}_i}^*(k|t)$, $\beta_i(0) > 0$ and $\sum_{i=1}^M \beta_i(0) \leq 1$. Using Proposition 1 in Mayne et al. (2005) it is possible to show that $\sum_{i=1}^M \beta_i(t) \leq 1$. For the negotiation parameter $\Delta\beta_i$ we introduce the local constraint (11c) following the distributed information structure.

Lemma 1. Let the neighborhood graph be undirected, i.e., if $i \in \mathcal{N}_j$ it follows that $j \in \mathcal{N}_i$. If the local constraints (11c) are enforced locally, then the global constraint $\sum_{i=1}^M \Delta\beta_i = 0$ is fulfilled.

Proof. We introduce an auxiliary matrix $M_a \in \mathbb{R}^{M \times M}$ defined as

$$[M_a]_{ij} = \begin{cases} \frac{\Delta\beta_j}{|\mathcal{N}_j| - 1} & \text{if } j \in \mathcal{N}_i \setminus i \\ 0 & \text{otherwise,} \end{cases} \quad (16)$$

where $|\mathcal{N}_j|$ denotes the number of subsystems in the set \mathcal{N}_j . The matrix M_a contains in each row i the terms $\frac{\Delta\beta_j}{|\mathcal{N}_j| - 1}$ for all $j \in \mathcal{N}_i \setminus i$ and in each column j only terms $\frac{\Delta\beta_j}{|\mathcal{N}_j| - 1}$. Because of the undirected graph structure, each column contains $|\mathcal{N}_j| - 1$ entries. If we multiply M_a by a vector of ones $\mathbf{1} \in \mathbb{R}^M$, we obtain a vector containing the local constraints

$$M_a \mathbf{1} = \begin{bmatrix} \sum_{j \in \mathcal{N}_1 \setminus 1} \frac{\Delta \beta_j}{|\mathcal{N}_j| - 1} \\ \vdots \\ \sum_{j \in \mathcal{N}_M \setminus M} \frac{\Delta \beta_j}{|\mathcal{N}_j| - 1} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}, \quad (17)$$

where the right-hand side is enforced locally. We can see that $\mathbf{1}^\top M_a \mathbf{1} = 0$. Otherwise, if we multiply the matrix M_a by $\mathbf{1}^\top$ from the left first, the $|\mathcal{N}_j| - 1$ entries in each column j sum up to $\Delta \beta_j$. This results in

$$\mathbf{1}^\top M_a \mathbf{1} = \sum_{i=1}^M \Delta \beta_i = 0, \quad (18)$$

which shows that enforcing the local constraints ensures the global constraint. \square

Remark 3.1. The local constraints provide a sufficient, but not a necessary condition. At the cost of being more conservative, (11c) is suited for distributed optimization. For example, for the circular graph structure in Figure 1(a) the constraint (11c) enforces all $\Delta \beta_i$ to be 0. For the centered graph structure in Figure 1(b) the constraint (11c) gives almost the same degree of freedom for choosing $\Delta \beta_i$ as the global constraint by only enforcing $\Delta \beta_1 = 0$ and $\sum_{i=2}^5 \Delta \beta_i = 0$.




Fig. 1. Examples of undirected neighborhood structures.

3.4 Structured Terminal Constraint

In robust MPC according to Mayne et al. (2005), the terminal nominal state is constrained to an invariant terminal set, i.e., $z(N|t) \in \mathcal{X}_f$, in order to ensure feasibility and stability. Since the real system state lies within a tube around the nominal state we can define the terminal safe set as $\mathcal{S}_f = \mathcal{X}_f \oplus \Omega$.

Following from the existence of a stabilizing feedback control law for system (4) according to Assumption 3.1, a structured ellipsoidal terminal invariant set $\mathcal{X}_f \subseteq \mathcal{X} \ominus \Omega$ with structured linear terminal feedback control law K_f can be computed using the approach in Conte et al. (2016) based on linear matrix inequalities. This allows to write a local terminal constraint for all $i \in \mathcal{M}$ as

$$z_i(N|t) \in \mathcal{X}_{f,i}(\alpha_i(t)) = \{z_i \in \mathbb{R}^{n_i} \mid z_i^\top P_{f,i} z_i \leq \alpha_i(t)\} \quad (19)$$

with the level set update rule

$$\alpha_i(t+1) = \alpha_i(t) + z_{\mathcal{N}_i}(t)^\top \Gamma_{\mathcal{N}_i} z_{\mathcal{N}_i}(t) \quad \forall i \in \mathcal{M} \quad (20)$$

with $\sum_{i=1}^M \alpha_i(0) \leq \alpha$ and $\alpha_i(0) \geq 0$. A synthesis procedure to obtain K_f , $P_{f,i}$, α and $\Gamma_{\mathcal{N}_i}$ is given in Conte et al. (2016).

3.5 Safety Guarantee

Based on the properties of distributed tube-based MPC according to Conte et al. (2013), problem (11) defines a recursively feasible distributed safety framework with (13) as safety control law.

Theorem 2. If Assumption 3.1 holds, then (13) is a safety control law with (12) as the corresponding safe set according to Definition 2.1. It also holds that (12) is an RPI set.

Proof. The proof follows the line of the proof of Theorem V.1 in Conte et al. (2013) and Theorem III.7 in Wabersich and Zeilinger (2018b). Assuming (11) is feasible at time step t with $x(t) \in \mathcal{X}_N$ we obtain an optimal solution \mathbf{z}^* , \mathbf{v}^* , \tilde{u}_i^* and $\Delta \beta_i^*$. We start by showing that the application of \tilde{u}_i^* results in $x(t+1) \in \mathcal{X}_N$, for which we can again find a feasible solution and ensure recursive feasibility using the properties of the terminal set.

As shown in Lemma 1, constraint (11c) ensures that $\sum_{i=1}^M \Delta \beta_i^* = 0$. Combined with $\sum_{i=1}^M \beta_i(t) \leq 1$ and (11d) this ensures that $\sum_{i=1}^M \tilde{\beta}_i(t) \leq 1$. Using (11b) and condition (10) it holds that $x(t) \in z^*(0|t) \oplus \Omega$. Following Proposition 1 in Mayne et al. (2005) the application of the tube-based control law (11i) with the nominal state dynamics (11f) results in $x(t+1) \in z^*(1|t) \oplus \Omega$. Because of condition (10) this can be written depending on local information as

$$\sum_{i=1}^M e_{\mathcal{N}_i}^*(1|t)^\top P_{\mathcal{N}_i} e_{\mathcal{N}_i}^*(1|t) \leq 1, \quad (21)$$

which by the definition of the update rule (15) shows that $\sum_{i=1}^M \beta_i(t+1) \leq 1$. The tightened nominal state constraints (11g) ensure $x(t+1) \in \mathcal{X}$.

Also by the definition of the update rule (15) we know that at time step $t+1$ the initial nominal state $z(0|t+1) = z^*(1|t)$ is a feasible solution ensuring that (11b) is fulfilled by choosing the negotiation parameters $\Delta \beta_i = 0$. Using the properties of the terminal set \mathcal{X}_f and control law $\sigma_f(z)$ a feasible solution at time step $t+1$ for all possible $x(t+1)$ is the series of nominal states $\{z^*(1|t), \dots, z^*(N|t), Az^*(N|t) + B\sigma_f(z^*(N|t))\}$, nominal inputs $\{v^*(1|t), \dots, v^*(N-1|t), \sigma_f(z^*(N|t))\}$ and the tube-based control law $(v^*(1|t) + K_\Omega(x(t+1) - z^*(1|t)))$. This shows by induction that the problem (11) is recursively feasible and (13) is a safety control law with (12) as the corresponding robust positively invariant safe set according to Definition 2.1. \square

Similarly to Corollary III.8 in Wabersich and Zeilinger (2018b) it is possible to state here that by the recursive feasibility of (11) a longer horizon N will lead to a larger or at least equally sized safe set (Wabersich and Zeilinger, 2018b).

Remark 3.2. In the case of no uncertainty in the system model, the backup control law reduces to a nominal MPC scheme according to Conte et al. (2016). Such a nominal safety framework reduces the complexity of practical implementations, while still allowing for well performing learning-based controllers for distributed systems. It also allows to handle systems with soft constraints.

4. DESIGN OF STRUCTURED TUBE

In this section we present a design procedure to obtain a structured ellipsoidal tube Ω based on Larsen et al. (2017) and Limon et al. (2008). The key advantage of the approach is that it allows to simultaneously compute a structured feedback control law K_Ω and optimize the shape of the tube for minimum constraint tightening. A related method has been introduced in Köhler (2017) with the difference that we consider ellipsoidal disturbance sets here. The synthesis problem introduced in Conte et al. (2013), in contrast, assumes that a structured feedback control law according to Assumption 3.1 is given. Obtaining such a control law following Conte et al. (2016) may, depending on the distributed system model, not lead to a tube fitting within the constraints even for small disturbance sets \mathcal{W}_i .

To simplify the notation, we define lifting matrices to switch from global to local state space, $x_i = T_i x$ with $T_i \in \{0, 1\}^{n_i \times n}$, from global to neighborhood state space, $x_{\mathcal{N}_i} = V_i x$ with $V_i \in \{0, 1\}^{n_{\mathcal{N}_i} \times n}$, and from global to local disturbances space, $w_i = W_i w$ with $W_i \in \{0, 1\}^{p_i \times p}$.

We define a structured ellipsoidal RPI set $\Omega^a = \{e \in \mathbb{R}^n \mid \sum_{i=1}^M e_i^\top P_i e_i \leq 1\}$ with all $P_i \in \mathbb{R}^{n_i \times n_i}$ positive semi-definite, i.e., P has a block-diagonal structure by definition, which can be obtained via the synthesis problem

$$\min_{E, K, h, o, \tau, \beta^s, \beta^{s+}} h^\top h + o^\top o \quad (22a)$$

$$\text{s.t. } \sum_{i=1}^M \beta_i^s = 1, \sum_{i=1}^M \beta_i^{s+} = 1 \quad (22b)$$

$\forall i \in \mathcal{M} :$

$$\begin{bmatrix} \tau_{M+1} \bar{E}_i & 0 & C_i^\top \\ 0 & \tau_i Q_i & G_i^\top \\ C_i & G_i & E_i \end{bmatrix} \succeq 0 \quad (22c)$$

$$\tau_{M+1} \beta_i^s - \beta_i^{s+} + \tau_i q_i \leq 0 \quad (22d)$$

$$\tau_i \geq 0, \tau_{M+1} \geq 0 \quad (22e)$$

$$\beta_i^s \geq 0, \beta_i^{s+} \geq 0 \quad (22f)$$

$\forall j \in \{1, \dots, n_{h_i}\} :$

$$\begin{bmatrix} [h_{\mathcal{N}_i}]_j^2 & [H_{\mathcal{N}_i}]_j E_{\mathcal{N}_i} \\ E_{\mathcal{N}_i} [H_{\mathcal{N}_i}]_j & E_{\mathcal{N}_i} \end{bmatrix} \succeq 0 \quad (22g)$$

$\forall j \in \{1, \dots, n_{o_i}\} :$

$$\begin{bmatrix} [o_i]_j^2 & [O_i K_i]_j \\ [K_i^\top O_i^\top]_j & E_{\mathcal{N}_i} \end{bmatrix} \succeq 0 \quad (22h)$$

where $E_i = P_i^{-1}$, $E_{\mathcal{N}_i} = P_{\mathcal{N}_i}^{-1}$, $P_{\mathcal{N}_i} = \text{diag}_{j \in \mathcal{N}_i}(P_j)$, $K_i = K_{\Omega, i} E_{\mathcal{N}_i}$, $\bar{E}_i = V_i T_i^\top E_i T_i V_i^\top \in \mathbb{R}^{n_{\mathcal{N}_i} \times n_{\mathcal{N}_i}}$ is lifted to the neighborhood space and $C_i = A_{\mathcal{N}_i} E_{\mathcal{N}_i} + B_i K_i$. $\tau_i \geq 0$ and $\tau_{M+1} \geq 0$ are S-procedure parameters. The constraints (22b)-(22h) are similar to (9b)-(9e) in Larsen et al. (2017), where constraints (22b)-(22f) encode the RPI implication (9) and (22g)-(22h) ensure that the resulting tube lies within the polytopic state and input constraints with right-hand side h and o . The result of the optimization problem (22) is the RPI set Ω^a , defined by the optimizers P_i^* and K_i^* , $i \in \mathcal{M}$. We can obtain the distributed stabilizing feedback control law $K_{\Omega, i}$ as $K_{\Omega, i} = K_i^* P_{\mathcal{N}_i}^*$.

The difference to Larsen et al. (2017) lies in the objective (22a), which is adapted from Limon et al. (2008) and in connection with the constraints (22g) and (22h) is chosen to minimize the tightening of the polytopic state and input constraints by the tube Ω^a . Note that the objective (22a) uses a unitary weighting under the assumption that the state and input constraints are given in a normalized form. Problem (22) thereby allows to optimize over the shape of Ω^a (defined by P) and the structured feedback control law ($K_{\Omega, i} = K_i P_{\mathcal{N}_i}$) at the same time.

Proposition 4.1. Consider P_i^* and K_i^* , $i \in \mathcal{M}$, resulting from problem (22). The structured ellipsoidal set $\Omega^a = \{e \in \mathbb{R}^n \mid \sum_{i=1}^M e_i^\top P_i^* e_i \leq 1\}$ is an RPI set for system (2) under the distributed linear control law $u_i(t) = K_{\Omega, i} x_{\mathcal{N}_i}(t)$ with $K_{\Omega, i} = K_i^* P_{\mathcal{N}_i}^*$.

Proof. The ellipsoidal set Ω^a is an RPI set according to Definition 3.1 if it fulfills (9), which we show in the line of the proof of Lemma 4 in Larsen et al. (2017). We introduce variables $\beta_i^s \geq 0$, $\beta_i^{s+} \geq 0$ with $\sum_{i=1}^M \beta_i^s = 1$, $\sum_{i=1}^M \beta_i^{s+} = 1$, to transform the corresponding terms in (9) to $\sum_{i=1}^M e_i(k|t)^\top P_i e_i(k|t) \leq \sum_{i=1}^M \beta_i^s$ and $\sum_{i=1}^M e_i(k+1|t)^\top P_i e_i(k+1|t) \leq \sum_{i=1}^M \beta_i^{s+}$. Using the S-procedure (see e.g., Boyd et al. (1994)) we can write (9) as the linear matrix inequality in (23), where $\bar{P}_i = T_i^\top P_i T_i$ and $\bar{Q}_i = W_i^\top Q_i W_i$ are lifted to global variable spaces, and $A_{cl} = (A + BK_\Omega)$. Due to the sparse structure of A_{cl} , we can write (23) as the two sufficient conditions in (24) for all $i \in \mathcal{M}$. Multiplying (24a) by $\text{diag}(P_{\mathcal{N}_i}^{-1}, I)$ from both, the left and right, and applying the Schur complement (see e.g., Boyd et al. (1994)) results in constraint (22c). Therefore, constraints (22b)-(22f) enforce the implication (9) and ensure that Ω^a is an RPI set. \square

Based on the solution P_i^* of problem (22), it is possible to introduce local sets $\Omega_i^a = \{e_i \in \mathbb{R}^{n_i} \mid e_i^\top P_i^* e_i \leq \beta_i^a\}$ such that $e(k|t) \in \Omega^a$ is equivalent to

$$\forall i \in \mathcal{M} \quad \exists \beta_i^a \geq 0 : e_i(k|t) \in \Omega_i^a, \sum_{i=1}^M \beta_i^a \leq 1. \quad (25)$$

Note that the existence of some local β_i^a for a given $e(k|t) \in \Omega^a$ is not related to the variables β_i^{s*} and β_i^{s+*} , which are introduced when designing (22) to derive distributed sufficient conditions for the global invariance condition (9).

Remark 4.1. While introducing β_i^s and β_i^{s+} in (22) can improve the synthesis of Ω^a , it requires the global consensus constraints in (22b). The additional design flexibility can, e.g., be beneficial when the scale of the local disturbances differs significantly, by allowing for a growing set for systems with larger disturbances, when other systems can compensate for the growth. Depending on the problem, β_i^s and β_i^{s+} can also be fixed to fully distribute (22), e.g., to $1/M$, while maintaining its feasibility.

Remark 4.2. The block-diagonal structure of Ω^a obtained through (22) is more restrictive compared to the block-sparse structure of the ellipsoidal tube introduced in Conte et al. (2013). However, since in (22) we simultaneously optimize over $K_{\Omega, i}$, one could also use this distributed controller to obtain a structured ellipsoidal tube using the synthesis problem (23) in Conte et al. (2013), with a feasible solution being $P_{\mathcal{N}_i} = V_i T_i^\top P_i^* T_i V_i^\top$ and $S_i = 0$.

$$\sum_{i=1}^M \left(\begin{bmatrix} A_{cl}^\top \bar{P}_i A_{cl} & A_{cl}^\top \bar{P}_i G & 0 \\ G^\top \bar{P}_i A_{cl} & G^\top \bar{P}_i G & 0 \\ 0 & 0 & -\beta_i^{s+} \end{bmatrix} + \tau_i \begin{bmatrix} 0 & 0 & 0 \\ 0 & -\bar{Q}_i & 0 \\ 0 & 0 & q_i \end{bmatrix} + \tau_{M+1} \begin{bmatrix} -\bar{P}_i & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \beta_i^s \end{bmatrix} \right) \preceq 0, \quad (23)$$

$$\begin{bmatrix} (A_{N_i} + B_i K_{\Omega,i})^\top P_i (A_{N_i} + B_i K_{\Omega,i}) - \tau_{M+1} V_i T_i^\top P_i T_i V_i^\top & (A_{N_i} + B_i K_{\Omega,i})^\top P_i G_i \\ G_i^\top P_i (A_{N_i} + B_i K_{\Omega,i}) & G_i^\top P_i G_i - \tau_i Q_i \end{bmatrix} \preceq 0, \quad (24a)$$

$$\tau_{M+1} \beta_i^s - \beta_i^{s+} + \tau_i q_i \leq 0. \quad (24b)$$




Fig. 2. Chain of M masses connected by springs and dampers.

5. NUMERICAL EXAMPLES

In this section we show the behavior of the introduced distributed safety framework in numerical simulations. All experiments were run using MATLAB on an Intel Core i7 2.7 GHz machine with 20 GB of RAM. We used the YALMIP toolbox Löfberg (2004) with MOSEK as solver. We show that the introduced safety certificate is able to ensure constraint satisfaction and compare the regulation of a distributed system to the origin using two different unsafe nominal controllers against directly applying robust distributed MPC according to Conte et al. (2013).

We consider a chain of M masses m_i interconnected by springs and dampers with local forces F_i as inputs to the system. A schematic of the system is shown in Figure 2.

Such a system could, e.g., be used to model a vehicle platoon, which possibly reduces energy consumption and enhances safety of road traffic (Li et al., 2015). In the following we consider a numerical example with $M = 9$ masses and parameter values $m_i = 1$, $k_{ij} = 0.1$ and $d_{ij} = 0.1$. The local disturbances are constrained to $\mathcal{W}_i = \{w_i \in \mathbb{R}^2 | w_i^\top w_i \leq 1.1e - 3\}$. We use a discrete-time model obtained by Euler forward method with sampling time of $0.2[s]$. To show constraint satisfaction we use a nominal feedback control law according to Conte et al. (2016) as input $u_{\mathcal{L}_i}$ to regulate the local subsystems to their local origins. This control law can be obtained from data by estimating the mean system dynamics. The local positions of the subsystems are constrained to $|p_i| \leq 1[m]$, with the exception of subsystem 2, where we set $-1[m] \leq p_2 \leq 0.1[m]$, the local velocities to $|v_i| \leq 1[m/s]$ and the inputs to $|u_i| \leq 5$. The local states are relative to a local coordinate system, which potentially is moving with a constant velocity. The introduced synthesis procedure in Section 4 with $\tau_{M+1} = 0.055$ and setting $\beta_i^s = \beta_i^{s+} = 1/M$ allowed to obtain a structured ellipsoidal tube contained within the state constraints. The synthesis procedure described in Conte et al. (2013) does not allow to adapt the shape of the tube to the tight constraint of subsystem 2, and did not result in a feasible tube to tighten the constraints for the same system parameters. We simulate the system for 20 time steps with a prediction horizon $N = 10$ for the safety framework. In Figure 3 we

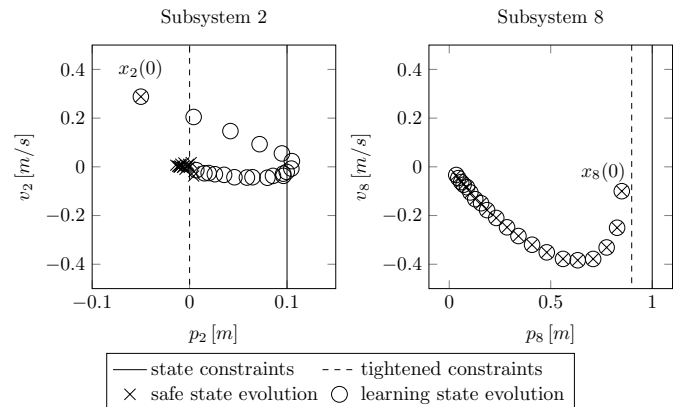


Fig. 3. Local state space (position and velocity) of the subsystems 2 and 8 in the chain with system evolution over 20 time steps.

can see that the distributed control law $u_{\mathcal{L}_2}$ slows down subsystem 2 very slowly and therefore leads to violation of the state constraints. The DMPSC framework is able to slow down the system on a short horizon and keep it safely around the origin. For subsystem 8 we can see, that the safety framework does not change the evolution of the local system and allows to directly apply $u_{\mathcal{L}_8}$.

For the same numerical system example, we analyzed the closed-loop performance according to local stage costs $l_i(x_{N_i}, u_i) = 1/2 \cdot x_{N_i}^\top I x_{N_i} + u_i^\top I u_i$ where I is an identity matrix with matching dimension. We compare two different well performing policies, where the safety framework is needed to ensure constraint satisfaction, with a robust distributed MPC scheme, which is able to handle state and input constraints. The three different controller variants are: The linear feedback control law combined with the safety certificate as in the example before (DMPSC 1), a nominal distributed MPC controller without terminal cost and constraints, combined with the safety certificate (DMPSC 2) and a robust distributed MPC control law according to Conte et al. (2013) (RDMPC). We use a prediction horizon of $N = 10$ for the DMPSC, the nominal and the robust MPC control law. The box-plot in Figure 4 shows the resulting closed-loop costs for 20 simulations with different initial positions over 20 simulation steps. We can see that the cases DMPSC 1 and 2 have lower costs than the case RDMPC. This shows that using the safety framework in connection with an unsafe controller can lead to a better performance. For our implementation, the case DMPSC 1 on median needed 35% less solver time compared to the case RDMPC. For the case DMPSC 2 the solver time was only about 7% lower, since finding the nominal MPC control law is more computationally intensive than a linear feedback control law.




Fig. 4. Box-plot of accumulated closed-loop costs.

6. CONCLUSION

In this paper we introduced a distributed safety framework, which is able to ensure constraint satisfaction of distributed linear systems with coupled dynamics and bounded additive disturbances in connection with any learning-based controller. In our numerical examples we showed, that this can potentially lead to a better performance in terms of cost, while being computationally more efficient.

REFERENCES

- Abbeel, P., Coates, A., and Ng, A.Y. (2010). Autonomous helicopter aerobatics through apprenticeship learning. *Int. J. Rob. Res.*, 29(13), 1608–1639.
- Akametalu, A.K., Fisac, J.F., Gillula, J.H., Kaynama, S., Zeilinger, M.N., and Tomlin, C.J. (2014). Reachability-based safe learning with gaussian processes. In *53rd IEEE Conf. on Decision and Control*, 1424–1431.
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P.F., Schulman, J., and Mané, D. (2016). Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.
- Berkenkamp, F., Turchetta, M., Schoellig, A., and Krause, A. (2017). Safe model-based reinforcement learning with stability guarantees. In *Advances in Neural Information Processing Systems 30*, 908–918. Curran Associates, Inc.
- Bertsekas, D.P. and Tsitsiklis, J.N. (1989). *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- Bhattacharai, B.P., Kouzelis, K., Mendaza, I.D.D.C., Bak-Jensen, B., Pillai, J.R., and Myers, K.S. (2017). Smart grid constraint violation management for balancing and regulating purposes. *IEEE Transactions on Industrial Informatics*, 13(6), 2864–2875.
- Boyd, S., El Ghaoui, L., Feron, E., and Balakrishnan, V. (1994). *Linear Matrix Inequalities in System and Control Theory*. SIAM studies in applied mathematics: 15.
- Buşoniu, L., Babuška, R., and De Schutter, B. (2010). *Multi-agent Reinforcement Learning: An Overview*, 183–221. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Conte, C., Zeilinger, M.N., Morari, M., and Jones, C.N. (2013). Robust distributed model predictive control of linear systems. In *2013 European Control Conf. (ECC)*, 2764–2769.
- Conte, C., Jones, C.N., Morari, M., and Zeilinger, M.N. (2016). Distributed synthesis and stability of cooperative distributed model predictive control for linear systems. *Automatica*, 69(C), 117–125.
- Fisac, J.F., Akametalu, A.K., Zeilinger, M.N., Kaynama, S., Gillula, J., and Tomlin, C.J. (2019). A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7), 2737–2752.
- García, J. and Fernández, F. (2015). A comprehensive survey on safe reinforcement learning. *J. Mach. Learn. Res.*, 16(1), 1437–1480.
- Hwangbo, J., Lee, J., Dosovitskiy, A., Bellicoso, D., Tsounis, V., Koltun, V., and Hutter, M. (2019). Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26).
- Köhler, J. (2017). *Distributed economic model predictive control under inexact minimization with application to power systems*. Master’s thesis, University of Stuttgart.
- Köhler, J., Müller, M.A., and Allgöwer, F. (2019). Distributed model predictive control - recursive feasibility under inexact dual optimization. *Automatica*, 102, 1–9.
- Langbort, C., Chandra, R.S., and D’Andrea, R. (2004). Distributed control design for systems interconnected over an arbitrary graph. *IEEE Transactions on Automatic Control*, 49(9), 1502–1519.
- Larsen, R.B., Carron, A., and Zeilinger, M.N. (2017). Safe learning for distributed systems with bounded uncertainties. *IFAC-PapersOnLine*, 50(1), 2536 – 2542. 20th IFAC World Congress.
- Li, S.E., Zheng, Y., Li, K., and Wang, J. (2015). An overview of vehicular platoon control under the four-component framework. In *2015 IEEE Intelligent Vehicles Symposium (IV)*, 286–291.
- Limon, D., Alvarado, I., Alamo, T., and Camacho, E. (2008). On the design of robust tube-based MPC for tracking. *IFAC Proceedings Volumes*, 41(2), 15333 – 15338. 17th IFAC World Congress.
- Löfberg, J. (2004). Yalmip : A toolbox for modeling and optimization in matlab. In *In Proc. of the CACSD Conf.* Taipei, Taiwan.
- Mayne, D.Q., Seron, M.M., and Raković, S.V. (2005). Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2), 219 – 224.
- Sartoretti, G., Wu, Y., Paivine, W., Kumar, T.K.S., Koenig, S., and Choset, H. (2019). Distributed reinforcement learning for multi-robot decentralized collective construction. In *Distributed Autonomous Robotic Systems*, 35–49. Springer International Publishing, Cham.
- Seto, D., Krogh, B., Sha, L., and Chutinan, A. (1998). The simplex architecture for safe online control system upgrades. In *Proc. of the 1998 American Control Conf. (ACC)*, volume 6, 3504–3508.
- Wabersich, K.P. and Zeilinger, M.N. (2018a). Scalable synthesis of safety certificates from data with application to learning-based control. In *2018 European Control Conf. (ECC)*, 1691–1697.
- Wabersich, K.P. and Zeilinger, M.N. (2018b). Linear model predictive safety certification for learning-based control. In *2018 IEEE Conf. on Decision and Control (CDC)*, 7130–7135.