Data-Driven Distributed Stochastic Model Predictive Control with Closed-Loop Chance Constraint Satisfaction

Simon Muntwiler*, Kim P. Wabersich*, Lukas Hewing, and Melanie N. Zeilinger

Abstract— Distributed model predictive control methods for uncertain systems often suffer from considerable conservatism and can tolerate only small uncertainties, due to the use of robust formulations that are amenable to distributed design and computational methods. In this work, we propose a distributed stochastic model predictive control (DSMPC) scheme for dynamically coupled linear discrete-time systems subject to unbounded additive disturbances that are potentially correlated in time. An indirect feedback formulation ensures recursive feasibility of the MPC problem, and a data-driven, distributed and optimization-free constraint tightening approach allows for exact satisfaction of chance constraints during closed-loop control, addressing typical sources of conservatism. The computational complexity of the proposed controller is similar to nominal distributed MPC. The approach is finally demonstrated in simulations for the temperature control of a large-scale data center subject to randomly varying computational loads.

I. INTRODUCTION

Sensing and communication capabilities are increasingly available in many technical systems, allowing interconnected systems to measure information locally and share it with other agents to optimize a common global objective. In the case of manufacturing systems, for instance, multiple machines may be necessary to assemble a product, and the use of each machine can be optimally scheduled based on shared information between the different production steps, increasing the overall efficiency. Solving such a large-scale control problem in a centralized manner, however, often results in intractable communication requirements or computationally infeasible optimization problems [1]. Distributed control algorithms address these issues by exploiting the distributed structure of the system and carry out computations locally while only requiring state information from neighboring subsystems.

In particular for large-scale systems, deriving an accurate system model and description of its operating conditions is a challenging task. In model predictive control (MPC), resulting uncertainties are often modeled through additive disturbances acting on the system. These can be addressed in a robust fashion, guaranteeing constraint satisfaction under *all* disturbance realizations in a compact set [2]. Distributed robust approaches, however, tend to introduce conservatism due to, for instance, enforcing a distributed structure on a robust positive invariant set [3] or handling dynamic couplings

as uncertainties [4]. A promising alternative is a stochastic approach, where an underlying stochastic characteristic of the disturbances is considered and constraints are satisfied with a certain probability level [5], providing a quantified assessment of risk. The closed-loop analysis for a stochastic MPC approach, however, is typically more challenging and often relies on bounded disturbance distributions [6], or unimodality and symmetric tightening assumptions [7], [8], [9]. Existing distributed stochastic MPC (DSMPC) frameworks assume Gaussian disturbance distributions [8] or general mean-variance information of i.i.d. disturbances [9]. In addition, chance constraints are usually enforced for all agents simultaneously, which can again introduce conservatism, in particular for large-scale systems.

Contributions: This paper introduces a distributed stochastic MPC (DSMPC) scheme for dynamically coupled linear systems and additive non-i.i.d. disturbances with potentially unbounded support. The considered goal is to regulate each local subsystem to its respective set-point, while satisfying local chance constraints with a given probability level, and thereby ensuring safety of each local subsystem. Instead of assuming a given distribution of the disturbance, we only assume to have access to samples of the disturbances either from experiments or simulation, resulting in a data-driven MPC formulation [10]. Making use of scenario optimization techniques [11], [12] and an indirect feedback formulation [13], we then provide guarantees for closed-loop chance constraint satisfaction for each individual agent.

In contrast to existing DSMPC approaches, recursive feasibility of the proposed DSMPC optimization problem is ensured by relying on indirect feedback as introduced in [13], where the actual measured state only enters the cost rather than the constraints. Chance constraints on states and inputs are handled using a data-based and distributed tightening approach, resulting in deterministic constraints on a nominal system state and input. In contrast to related schemes (e.g., [8]), the constraint tightening does not introduce additional conservatism compared to a centralized solution and allows to handle the local chance constraints in a non-conservative manner. The resulting optimization problem with respect to nominal states and inputs has computational complexity comparable to a nominal distributed MPC problem.

Related Work: DSMPC algorithms based on distributional information of the disturbances have been introduced for linear systems with zero-mean i.i.d. additive Gaussian disturbances in [8], and extended to output-feedback in [9]. These techniques similarly consider unbounded disturbance distributions, but cannot ensure recursive feasibility of the

The authors are with the Institute for Dynamic Systems and Control, ETH Zürich, Zürich, Switzerland {`simonmu,wkim,lhewing, mzeilinger`}`@ethz.ch`.

MPC problem directly; instead they make use of a recovery initialization in case of infeasibility. As a result, constraint satisfaction is guaranteed under symmetric tightening and unimodal disturbances only. For bounded disturbances and dynamically decoupled systems with coupling constraints, approaches ensuring constraint satisfaction were presented in [14]. A DSMPC framework based on disturbance samples, instead of distributional information, for linear systems with parameter uncertainty and additive disturbances was investigated in [15]. The constraints in the online DSMPC problem, however, need to be fulfilled for the entire set of disturbance samples, potentially resulting in a large number of constraints. Recursive feasibility and stability were not investigated.

Structure: We begin by introducing the problem formulation in Section II. In Section III, we derive the proposed DSMPC framework and show recursive feasibility. In Section IV, we specify the data-driven and distributed approach for chance constraint tightening and prove chance constraint satisfaction of the resulting DSMPC framework. Section V demonstrates the results for a numerical simulation example and Section VI concludes the paper.

II. PRELIMINARIES

A. Notation

A stacked vector $v \in \mathbb{R}^n$ consisting of subvectors $v_i \in \mathbb{R}^{n_i}$ with $i \in \mathcal{M} \subseteq \mathbb{N}$ is denoted as $v = \mathrm{col}_{i \in \mathcal{M}}(v_i)$. The distribution $\mathcal{Q}$ of a random variable w is denoted as $w \sim \mathcal{Q}$ and probabilities and conditional probabilities as $\mathrm{Pr}(A)$ and $\mathrm{Pr}(A|B)$ respectively. By $\mathbb{E}_w(x)$ we denote the expected value of x w.r.t. the random variable w.

B. Problem Formulation

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

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

with local state $x_i(t) \in \mathbb{R}^{n_i}$, input $u_i(t) \in \mathbb{R}^{m_i}$ and stochastic disturbance $w_i(t) \in \mathbb{R}^{p_i}$ for each subsystem i at time step t, where $A_{ij} \in \mathbb{R}^{n_i \times n_j}$, $B_i \in \mathbb{R}^{n_i \times m_i}$, and $G_i \in \mathbb{R}^{n_i \times p_i}$. We denote the set of indices of all subsystems as $\mathcal{M} = \{1, \ldots, M\}$. The set of neighbors $\mathcal{N}_i$ of subsystem i contains all indices of subsystems j, for which A_{ij} includes nonzero entries. We assume that each subsystem is able to exchange information with all other subsystems in its neighborhood. 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 $x_{\mathcal{N}_i}(t) = \mathrm{col}_{j \in \mathcal{N}_i}(x_j(t)) \in \mathbb{R}^{n_{\mathcal{N}_i}}$. Each subsystem i is subject to n_i^x half-space chance constraints on the local states and n_i^u half-space chance constraints on the local inputs

$$\mathrm{Pr}(h_{i,j}^{x\top} x_i(t) \leq 1) \geq p_{i,j}^x, \ j \in \{1, \ldots, n_i^x\} \quad (3a)$$

$$\mathrm{Pr}(h_{i,j}^{u\top} u_i(t) \leq 1) \geq p_{i,j}^u, \ j \in \{1, \ldots, n_i^u\} \quad (3b)$$

where $h_{i,j}^x \in \mathbb{R}^{n_i}$, $h_{i,j}^u \in \mathbb{R}^{m_i}$ and the probabilities are understood conditioned on the initial state.

The objective is to control the distributed stochastic system over a potentially large, but finite, task horizon $\bar{N}$ while satisfying the chance constraints (3) at every time step t. The stochastic disturbance sequence over the task horizon is assumed to be distributed according to $W = \left[\mathrm{col}_{i \in \mathcal{M}}(w_i(0))^\top, \ldots, \mathrm{col}_{i \in \mathcal{M}}(w_i(\bar{N}))^\top \right]^\top \sim \mathcal{Q}$, which can be a non-i.i.d. and correlated disturbance sequence with unbounded support. It is not necessary to know the distribution of the disturbances, but we assume to have access to samples from the distribution over the entire task horizon. Handling unbounded disturbances is especially important when the distribution and possibly existing bounds are not known in advance, with normal distributions as important special case.

In this paper, we introduce a distributed stochastic MPC scheme to approximate the solution of the optimal stochastic control problem by solving a simplified problem over a shorter horizon $N \ll \bar{N}$ in a receding horizon fashion. The local system dynamics (2) are split into a nominal state $z_i(t)$ and error $e_i(t)$ such that $x_i(t) = z_i(t) + e_i(t)$, as well as a nominal input $v_i(t)$ and potentially nonlinear tube controller $\pi_i(e_{\mathcal{N}_i}(t))$, resulting in

$$z_i(t+1) = A_{\mathcal{N}_i} z_{\mathcal{N}_i}(t) + B_i v_i(t), \quad (4a)$$

$$e_i(t+1) = A_{\mathcal{N}_i} e_{\mathcal{N}_i}(t) + B_i \pi_i(e_{\mathcal{N}_i}(t)) + G_i w_i(t), \quad (4b)$$

$$x_i(t) = z_i(t) + e_i(t), \quad (4c)$$

$$u_i(t) = v_i(t) + \pi_i(e_{\mathcal{N}_i}(t)), \quad (4d)$$

with initial condition $z_i(0) = x_i(0)$, and therefore $e_i(0) = 0$, and $z_{\mathcal{N}_i}(t) = \mathrm{col}_{j \in \mathcal{N}_i}(z_j(t)) \in \mathbb{R}^{n_{\mathcal{N}_i}}$ and $e_{\mathcal{N}_i}(t) = \mathrm{col}_{j \in \mathcal{N}_i}(e_j(t)) \in \mathbb{R}^{n_{\mathcal{N}_i}}$. The MPC problem optimizes the nominal input $v_i(t)$, while the tube controller $\pi_i(e_{\mathcal{N}_i}(t))$ is used to regulate deviations from the nominally planned trajectory. A feedback controller with such properties can be obtained in a distributed manner using, e.g., the methods in [16].

Remark 1: The considered class of tube controllers π_i includes controllers with saturation, allowing for the treatment of hard input constraints (e.g., due to physical actuator limits, see also [12]).

In the following, we introduce the recursively feasible distributed stochastic MPC scheme based on an indirect feedback formulation [13] in Section III. In Section IV, we then detail the tightening of nominal system states and inputs, which is performed in an optimization-free and distributed manner, ensuring chance constraint satisfaction in closed-loop operation.

III. DISTRIBUTED STOCHASTIC MODEL PREDICTIVE CONTROL

We aim to solve the stochastic control task over the task length $\bar{N}$ by employing a receding horizon control formulation over a shortened horizon N, i.e., the problem is repeatedly solved at each time step based on the current

measured state. Using the separation into nominal and error system in (4), the proposed DSMPC problem is given by

$$\min_{\mathbf{v}} \sum_{i=1}^{M} \mathbb{E}_{W_i^t} \left(l_f(x_i(N|t)) + \sum_{k=0}^{N-1} l_{t+k}(x_i(k|t), u_i(k|t)) \right) \tag{5a}$$

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

$$x_i(0|t) = x_i(t),\ z_i(0|t) = z_i(1|t-1),\ e_i(0|t) = e_i(t) \tag{5b}$$

$$z_i(N|t) = 0 \tag{5c}$$

$$W_i(t) = \left[w_i(0|t)^\top, \ldots, w_i(N|t)^\top \right] \sim \mathcal{Q}_i(t) \tag{5d}$$

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

$$z_i(k+1|t) = A_{\mathcal{N}_i} z_{\mathcal{N}_i}(k|t) + B_i v_i(k|t) \tag{5e}$$

$$x_i(k+1|t) = z_i(k+1|t) + e_i(k+1|t) \tag{5f}$$

$$e_i(k+1|t) = A_{\mathcal{N}_i} e_{\mathcal{N}_i}(k|t)$$
$$\qquad + B_i \pi_i(e_{\mathcal{N}_i}(k|t)) + G_i w_i(k|t) \tag{5g}$$

$$u_i(k|t) = v_i(k|t) + \pi_i(e_{\mathcal{N}_i}(k|t)) \tag{5h}$$

$$h_{i,j}^{x\top} z_i(k|t) \le 1 - c_{i,j,t+k}^x \forall j \in \{1, \ldots, n_i^x\} \tag{5i}$$

$$h_{i,j}^{u\top} v_i(k|t) \le 1 - c_{i,j,t+k}^u \forall j \in \{1, \ldots, n_i^u\} \tag{5j}$$

where $\mathbf{v} = \mathrm{col}_{k \in \{0,\ldots,N-1\}}(v(k|t))$ with $v(k|t) = \mathrm{col}_{i \in \mathcal{M}}(v_i(k|t))$. For $k \in \{0, \ldots, N\}$ and every subsystem i, the vector $x_i(k|t) \in \mathbb{R}^{n_i}$ denotes the k-steps ahead predicted state computed at time step t, and $z_i(k|t)$, $v_i(k|t)$, $e_i(k|t)$ and $u_i(k|t)$ the predicted nominal state, nominal input, error and input, respectively. The input (4d) applied to system (1) is then defined by the solution $\mathbf{v}^*$ as

$$v_i(t) = v_i^*(0|t). \tag{6}$$

Most commonly, robust and stochastic MPC schemes initialize the *nominal* state $z_i(0|t)$ with the current measured state at time step t, i.e., $x_i(t)$. In the stochastic setting, this can easily lead to feasibility issues, in particular due to the potentially unbounded nature of the stochastic disturbance [7]. In an indirect feedback stochastic MPC formulation [13], the nominal state is instead initialized at each time step t with the first predicted nominal state $z_i(1|t-1)$ obtained at time step $t-1$, while the state measurement $x_i(t)$ initializes $x_i(0|t)$ (see (5b)). Note that via the optimization of the objective (5a) with respect to $x_i(k|t)$, feedback is also introduced on the nominal state evolution $z_i(t)$, hence it is referred to as indirect feedback.

As a result of this initialization, the nominal dynamics in (4a) are valid in closed-loop operation. Note that this is not the case, if $z_i(0|t)$ is optimized, as often the case in robust tube MPC formulations [17], or if it is set equal to the measured state $x_i(t)$. From this nominal state evolution, it follows that the closed-loop error evolves independently of the MPC optimization according to (4b) and can therefore be simulated forward by only having access to samples of the disturbances $w_i(t)$. This allows to precompute the error prediction (5g) prior to solving the optimization problem (5).

The expectation in the objective (5a) is taken with respect to a disturbance sequence $W_i(t)$ over the prediction horizon distributed according to $\mathcal{Q}_i(t)$ (5d). For disturbances correlated in time, this addresses the fact that past disturbances provide information which can be utilized in the optimization of the cost. This can be achieved, e.g., by considering the marginal disturbances for each agent conditioned on past disturbance realizations $p(W_i(t)) = p\left(W_i(t)|\left[w_{\mathcal{N}_i}(0)^\top, \ldots, w_{\mathcal{N}_i}(t-1)^\top\right]^\top\right)$.

The expectation can be evaluated for the special case of i.i.d. disturbances and quadratic costs by considering only the mean of the predicted state and input, see [13]. For a general cost, it is not possible to analytically evaluate the expectation in (5a), but it can be approximated based on $N_{s,i}^{MPC}$ samples of the disturbance sequence $W_i(t)$ over the prediction horizon N for each subsystem i. The number of samples trade off prediction accuracy against online computational complexity.

Problem (5) makes use of tightened constraints on the nominal state and input of each subsystem in (5i) and (5j) to realize the chance constraints in (3). While the local error feedback π_i aims at reducing deviations from the nominally planned trajectory $z_i(t)$, the unknown disturbances $w_i(t)$ cause a non-vanishing error $e_i(t)$ for all $t \ge 0$, which can cause closed-loop constraint violations, even if $h_{i,j}^{x\top} z_i(t) \le 1$ and $h_{i,j}^{u\top} v_i(t) \le 1$ holds. Similar to ideas from robust MPC, we therefore introduce tightened half-space constraints using suitable tightening values $c_{i,j,t}^x$ and $c_{i,j,t}^u$. In Section IV, we introduce a data-driven and distributed method to compute $c_{i,j,t}^x$ and $c_{i,j,t}^u$ depending on the distribution of the trajectories of the error system (4b), such that the chance constraints (3) are fulfilled non-conservatively with the desired probability level.

The optimization problem in (5) can be solved in a distributed manner using distributed optimization techniques, see e.g., [18], [19], since the objective and constraints are only coupled between neighboring subsystems. This results in a fully distributed offline and online procedure.

Remark 2: For simplicity, we use a terminal equality constraint in (5). A less restrictive terminal constraint as similarly proposed in [12] could be integrated by using a distributed robust positive invariant terminal set, e.g., based on the results introduced in [3].

Recursive feasibility of the distributed MPC scheme (5) can be directly established using results from standard nominal MPC, because the stochastic variables only affect the objective of problem (5) and the constraints make use of the nominal system.

Theorem 1: If the optimization problem (5) is feasible for $x_i(0) = z_i(0)$, then applying the distributed control input (4d) with (6) to the dynamic system (1), problem (5) is feasible for all time steps $0 \le t \le \bar{N} - N$.

Proof: The distributed constraints in (5) can be combined to the centralized formulation and the proof follows the standard argument in MPC using the shifted sequence from the previous time step, as similarly shown in [12] for

Algorithm 1 Computation of tightening values for all subsystems i, time steps t, and half spaces j.

Input: Chance constraints (3), confidence level β, $l = 1, 2, .., N_s$ samples $W^{(l)}$.

Output: Tightening values $c_{i,j,t}^x$ and $c_{i,j,t}^u$.

1: **for** every sample $l = 1, 2, .., N_s$ **do**
2:    $(e_i^{(l)}(t), \pi_i^{(l)}(t)) \leftarrow$ distributed simulation of error system (4b) and corresponding feedback using disturbances $W^{(l)}$ and initial condition $e_i^{(l)}(0) = 0$.
3: **end for**
4: **for** every agent $i = 1, .., M$, time step $t = 0, .., \bar{N}$ **do**
5:    **for** every half-space $j = 1, .., n_i^x$ **do**
6:       Compute $c_{i,j,t}^x$ via Alg. 2 and $h_{i,j}^x, \{e_i^{(l)}(t)\}, p_{i,j}^x, \beta$
7:    **end for**
8:    **for** every half-space $j = 1, .., n_i^u$ **do**
9:       Compute $c_{i,j,t}^u$ via Alg. 2 and $h_{i,j}^u, \{e_i^{(l)}(t)\}, p_{i,j}^u, \beta$
10:    **end for**
11: **end for**

Algorithm 2 Single half-space tightening computation.

Input: Half-space direction $h \in \mathbb{R}^q$, samples $\xi^{(l)} \in \mathbb{R}^q$ with $l = 1, .., N_s$, probability level p, and confidence $1 - \beta$
Output: Tightening value c

1: $N_d \leftarrow (1-p)N_s - \sqrt{2(1-p)N_s \ln\left(\frac{1}{\beta}\right)}$
2: **while** number of $\{\xi^{(l)}\} > N_s - N_d$ **do**
3:    discard $\xi^{(l^*)}(t)$ with $l^* \leftarrow \arg\max_l h^\top \xi^{(l)}$
4: **end while**
5: $c \leftarrow \max_l h^\top \xi^{(l)}$

---

the centralized case. ∎

Remark 3: For the special case of a quadratic stage cost, Gaussian disturbances and a terminal weight satisfying the Lyapunov equation, the asymptotic convergence property shown in [13] can be extended to the distributed case.

IV. Distributed Data-driven constraint tightening

In the following, we derive a distributed, data-driven and optimization-free algorithm based on scenario optimization as proposed in [12], to obtain the tightening of nominal state and input constraints in (5i) and (5j). The proposed tightening allows to satisfy the chance constraints (3) in a non-conservative manner, meaning that if a constraint on the nominal state $z_i(t)$ is active, the probability of the real state $x_i(t)$ and applied input $u_i(t)$ violating the constraints is exactly $1-p_{i,j}^x$ and $1-p_{i,j}^u$ as specified in (3). Compared to related robust approaches, such as [3], the proposed design procedure avoids the solution of a distributed optimization problem involving bilinear matrix inequalities to determine the constraint tightening, by making use of samples of the closed-loop error according to the dynamics (4b).

Specifically, we compute tightening values $c_{i,j,t}^x$ and $c_{i,j,t}^u$, which ensure that the real local state $x_i(t) = z_i(t) + e_i(t)$ and input $u_i(t) = v_i(t) + \pi_i(e_{\mathcal{N}_i}(t))$ satisfy the half-space chance

constraints (3) at the desired probability level if the tightened nominal constraints (5i) and (5j) are active. Therefore, we choose $c_{i,j,t}^x$ and $c_{i,j,t}^u$ such that for each constraint j, time step t and subsystem i

$$\Pr(h_{i,j}^{x\top} e_i(t) \leq c_{i,j,t}^x) \geq p_{i,j}^x, \qquad j \in \{1, \dots, n_i^x\}$$
$$\Pr(h_{i,j}^{u\top} \pi_i(e_{\mathcal{N}_i}(t)) \leq c_{i,j,t}^u) \geq p_{i,j}^u, \quad j \in \{1, \dots, n_i^u\},$$

holds, bounding the distribution of the local error dynamics (4b) and error feedback in the local half-space directions $h_{i,j}^x$ and $h_{i,j}^u$. The tightening values $c_{i,j,t}^x$ and $c_{i,j,t}^u$ can be obtained by solving the stochastic optimization problem

$$c_{i,j,t}^x = \min c_x \text{ s.t. } \Pr(h_{i,j}^{x\top} e_i(t) \leq c_x) \geq p_{i,j}^x, \qquad (7a)$$
$$c_{i,j,t}^u = \min c_u \text{ s.t. } \Pr(h_{i,j}^{u\top} \pi_i(e_{\mathcal{N}_i}(t)) \leq c_u) \geq p_{i,j}^u. \quad (7b)$$

Relying on N_s samples of the disturbance sequences $W^{(l)}$ allows us to use arguments from scenario-based optimization (see e.g., [20], [11]) to approximate (7) by enforcing $h_{i,j}^{x\top} e_i^{(l)}(t) \leq c_x$ and $h_{i,j}^{u\top} \pi_i(e_{\mathcal{N}_i}^{(l)}(t)) \leq c_u$ as deterministic constraints for sampled error trajectories $e_i^{(l)}(t)$ based on disturbance samples $W^{(l)}$ as detailed in Algorithm 1. In fact, scenario-based optimization arguments [11] provide a confidence level $1 - \beta$ at which the sample-based solution fulfills the probabilistic constraints in (7) and even allows to discard a certain fraction of the most restrictive samples. The procedure is outlined in Algorithm 1 and 2.

Algorithm 1 takes the chance-constraints (3) as inputs, as well as disturbance samples $W^{(l)}$, and the confidence level parameter β, where $1-\beta$ corresponds to the confidence level of the scenario optimization, i.e. the confidence at which the computed constraint tightening results in closed-loop chance constraint satisfaction. In a first step, we generate the relevant error scenarios by simulating the error system for each disturbance sample, see Algorithm 1, lines 1-3. Note that the disturbance samples $W^{(l)}$ can be stored distributedly and that the simulation is a distributed operation requiring only neighbor-to-neighbor communication and therefore scales to arbitrarily large networks. After generating the error scenarios, every agent can approximately solve (7) for each state and input half-space separately in lines 5-10 using the subroutine in Algorithm 2. In Algorithm 2, line 1, we first compute the number of scenarios N_d that can be discarded based on the desired probability level p and confidence level $1-\beta$, see [12], [11] for details. To determine the required half-space level c, we iterate over the disturbance samples and discard the N_d most restricting samples in Algorithm 2, lines 2-4. The most restrictive remaining disturbance sample is then used to obtain the required tightening value in line 5. Note that increasing the number of samples N_s either allows to achieve a higher probability level p for the chance constraints, or a higher confidence level $1 - \beta$ of the scenario optimization problem. Since the required number of samples scales logarithmically with β, the confidence level can typically be chosen very high [11]. Note that the chosen number of samples $N_{s,i}^{MPC}$ to approximate the MPC cost is not related to the number of samples N_s to perform the constraint tightening and does not

affect constraint satisfaction guarantees. In fact, one would typically have $N_s \gg N_{s,i}^{MPC}$ since the number of samples for constraint tightening does not affect the online computation, and the required offline computations are reasonably cheap.

Recursive feasibility of problem (5) as shown in Theorem 1 and the tightened constraints on the nominal states (5i) and inputs (5j) with constants $c_{i,j,t}^x$ and $c_{i,j,t}^u$ obtained using Algorithm 1 allow to establish a guarantee for the satisfaction of the chance constraints (3) on states $x_i(t)$ and inputs $u_i(t)$ of each subsystem in closed-loop.

Theorem 2: Let $c_{i,j,t}^x$ and $c_{i,j,t}^u$ be obtained using Algorithm 1 and the control law (4d) with (6) be applied to the distributed system (1). With probability $1 - \beta$, the resulting local states $x_i(t)$ and inputs $u_i(t)$ satisfy the chance constraints in (3).

Proof: The proof follows the proof of Theorem 3 in [12], which is summarized here for completeness. Algorithm 1 greedily discards N_d of the initial N_s samples $e_i^{(l)}$ and sets the tightening value $c_{i,j,t}^x$ as the maximum over the remaining samples of $h_{i,j}^x e_i^{(l)}(t)$ via Algorithm 2 line 5. Therefore, for all remaining samples it holds that $h_{i,j}^x e_i^{(l)}(t) \le c_{i,j,t}^x$. From scenario optimization, we then have with probability $1 - \beta$, that $\Pr(h_{i,j}^{x\top} e_i(t) \le c_{i,j,t}^x) \ge p_{i,j}^x$. Therefore, constraining the local nominal state $z_i(t)$ to the tightened constraints (5i) results in the real state of the system $x_i(t) = z_i(t) + e_i(t)$ fulfilling the chance constraints in (3). The same arguments hold for the input constraints by using Algorithm 1 to obtain the tightening values $c_{i,j,t}^u$. ∎

Remark 4: The constraint satisfaction property in Theorem 2 renders the proposed DSMPC framework suitable for safety certification of distributed learning-based controllers in the line of [21], i.e., using a distributed MPC to verify and modify a proposed learning input if necessary. While satisfaction of constraints can only be ensured in probability, the computational complexity and conservatism can be dramatically reduced compared with other distributed safety verification schemes [21], [22].

Remark 5: For disturbances with zero mean and known variance, e.g., $W \sim \mathcal{N}(0, \Sigma_W)$ and a distributed linear tube control law $\pi_i(e_{\mathcal{N}_i}(t)) = K_i e_{\mathcal{N}_i}(t)$ with $K_i \in \mathbb{R}^{n_i \times n_{\mathcal{N}_i}}$, one can analytically compute mean and variance of the error sequence. Instead of a data-based tightening, an analytic tightening is then possible using the marginal local and neighborhood variances $\Sigma_i^e(t)$ and $\Sigma_{\mathcal{N}_i}^e(t)$, e.g., as

$$c_{i,j,t}^x = \phi^{-1}(p_{i,j}^x)\sqrt{h^\top \Sigma_i^e(t) h}, \qquad j \in \{1, \dots, n_i^x\},$$
$$c_{i,j,t}^u = \phi^{-1}(p_{i,j}^u)\sqrt{h^\top K_i \Sigma_{\mathcal{N}_i}^e(t) K_i^\top h}, \quad j \in \{1, \dots, n_i^u\},$$

where ϕ^{-1} is the quantile function of the standard normal distribution and all computations can be easily carried out in a distributed manner. A related approach computing the full variance matrix was presented in [8], where, using a possibly conservative additional step, guarantees are given for all subsystems simultaneously.

V. SIMULATION EXAMPLE

To highlight the effect of the proposed DSMPC scheme we consider the example of a distributed cooling system as used in [23]. The task of the cooling system is to control the temperature of a server farm, which can similarly be interpreted, e.g., as the temperature of production machines in a big manufacturing plant. Each local subsystem thereby has a heat source (e.g., heat production due to the computational load) and a cooling component (e.g., a fan or water cooling system). The temperature of each subsystem is affecting the temperature of neighboring systems. Cooling of the system is important in order to prevent defects due to overly high temperatures or safety shutdowns. At the same time, excessive cooling should be prevented.

We consider a server farm with $M = 100$ servers which on average have thermal couplings with 22.4 neighboring systems. Each server is heating up due to running computations and influences the servers in a circle of radius $r_{\max}$ in its proximity. Disturbances acting on each local server mimic the temperature increase (or decrease) due to high (or low) computations compared to the average computational load acting on the servers. The computational load is assumed to have a known time-varying mean over the course of the day.

The local system dynamics are defined as

$$x_i(t+1) = 1.01 x_i(t) + \sum_{j \in \mathcal{N}_i \setminus i} \frac{0.01}{1 + r_{ij}} x_j(t) + u_i(t) + w_i(t),$$
$$(8)$$

where $x_i(t)$ denotes the deviation from a desired temperature of operation $\bar{T}_i = 25°C$, with the actual temperature $T_i(t) = \bar{T}_i + x_i(t)$, $u_i(t)$ denotes a local cooling input, $w_i(t)$ the disturbance acting on each server and $0 \le r_{ij} \le r_{max}$ the distance between server i and j. We introduce state and input constraints as

$$-5 \le x_i \le 5 \quad \forall i \in \mathcal{M} \tag{9a}$$
$$-1 \le u_i \le 1 \quad \forall i \in \mathcal{M}. \tag{9b}$$

with desired probability level of 0.9. We use a tube controller $\pi_i(e_i(t)) = -0.5 e_i(t)$ for every subsystem resulting in the closed-loop nominal error system

$$e_i(t+1) = 0.51 e_i(t) + \sum_{j \in \mathcal{N}_i \setminus i} \frac{0.01}{1 + r_{ij}} e_j(t) \tag{10}$$

which is stable according to the Gersgorin Disk Theorem [24] if for all subsystems i

$$0.51 + \sum_{j \in \mathcal{N}_i \setminus i} \frac{0.01}{1 + r_{ij}} < 1. \tag{11}$$

We simulate the behavior of the DSMPC scheme for the temperature control system with a sampling time of $0.5h$ and a prediction horizon of $N = 12h$ over an effective task horizon $\bar{N} - N = 2d$. For the constraint tightening we sample $N_s = 100$ disturbance sequences that are correlated in time over the task horizon for each subsystem. As an example, disturbance samples for subsystem 9 and the actual
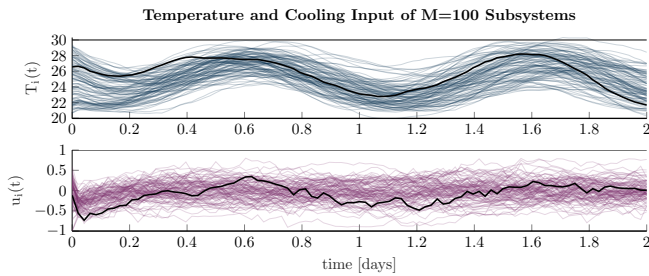
Fig. 1. Evolution of the temperature and control input for all subsystems over time with the one of subsystem 9 indicated in black.
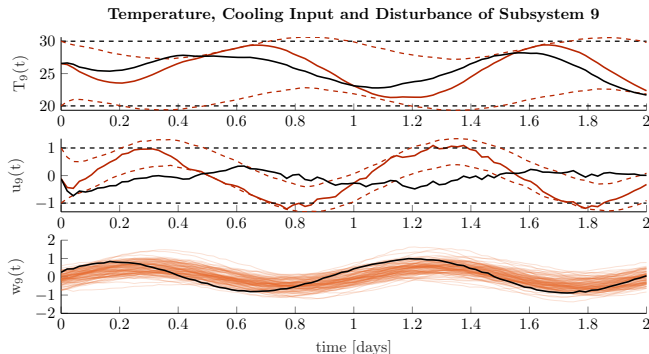

Fig. 2. The upper two subplots show the evolution of the real temperature and input (in black), nominal temparature and input (in red) and real and tightened constraints (dashed lines) for subsystem 9. The third subplot shows the disturbance samples used for tightening the constraints (in orange) and the actual disturbance acting on the system (in black).

disturbance are shown in the third subplot of Figure 2. The local stage costs are assumed to have the form

$$l_t(x_i, u_i) = x_i^\top x_i + 1000 u_i^\top u_i \tag{12}$$

which represents high cooling costs and the MPC cost is approximated using $N_{s,i}^{MPC} = 10$ samples for each subsystem.

Figure 1 shows the temperature evolution with the corresponding inputs over the course of two days. We can see that the chance constraints on the states are violated for only two subsystems, while the input constraints always hold. The upper two subplots of Figure 2 shows the temperature and input of subsystem 9 including the nominal states and inputs and the time-varying tightened constraints on nominal state and input. It shows that the constraints on the nominal state and input is active at several instances in time.

VI. CONCLUSIONS

In this paper, we introduced a DSMPC framework that ensures recursive feasibility, based on an indirect feedback formulation, and satisfaction of chance constraints in closed-loop in a non-conservative manner due to a data-driven and optimization-free constraint tightening approach. Both, the offline controller synthesis as well as the online operation can be performed in a completely distributed manner, offering a scalable and high performance DSMPC scheme with reduced conservatism compared with the literature.

REFERENCES

[1] C. Langbort, R. S. Chandra, and R. D'Andrea, "Distributed control design for systems interconnected over an arbitrary graph," *IEEE Trans. on Automatic Control*, vol. 49, no. 9, pp. 1502–1519, 2004.

[2] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in identification and control*. Springer, 1999, pp. 207–226.

[3] C. Conte, M. N. Zeilinger, M. Morari, and C. N. Jones, "Robust distributed model predictive control of linear systems," in *2013 European Control Conference (ECC)*, 2013, pp. 2764–2769.

[4] Dong Jia and B. Krogh, "Min-max feedback model predictive control for distributed control with communication," in *Proc. of the 2002 American Control Conference*, vol. 6, 2002, pp. 4507–4512.

[5] A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research," *IEEE Control Systems Magazine*, vol. 36, no. 6, pp. 30–44, 2016.

[6] M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgwer, "Constraint-tightening and stability in stochastic model predictive control," *IEEE Trans. on Automatic Control*, vol. 62, no. 7, pp. 3165–3177, 2017.

[7] L. Hewing and M. N. Zeilinger, "Stochastic model predictive control for linear systems using probabilistic reachable sets," in *2018 IEEE Conference on Decision and Control (CDC)*, 2018, pp. 5182–5188.

[8] C. Mark and S. Liu, "Distributed stochastic model predictive control for dynamically coupled linear systems using probabilistic reachable sets," in *Proc. 18th European Control Conf. (ECC)*, 2019, pp. 1362–1367.

[9] ——, "A stochastic output-feedback MPC scheme for distributed systems," *arXiv preprint arXiv:2001.10838*, 2020.

[10] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, "Learning-based model predictive control: Toward safe learning in control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, 2020.

[11] M. C. Campi and S. Garatti, "A sampling-and-discarding approach to chance-constrained optimization: feasibility and optimality," *Journal of Optimization Theory and Applications*, vol. 148, no. 2, pp. 257–280, 2011.

[12] L. Hewing and M. N. Zeilinger, "Scenario-based probabilistic reachable sets for recursively feasible stochastic model predictive control," *IEEE Control Systems Letters*, vol. 4, no. 2, pp. 450–455, 2020.

[13] L. Hewing, K. P. Wabersich, and M. N. Zeilinger, "Recursively feasible stochastic model predictive control using indirect feedback," *arXiv preprint arXiv:1812.06860*, 2018.

[14] L. Dai, Y. Xia, Y. Gao, and M. Cannon, "Distributed stochastic MPC of linear systems with parameter uncertainty and disturbances," in *Proc. 35th Chinese Control Conf. (CCC)*, 2016, pp. 4312–4317.

[15] V. Rostampour and T. Keviczky, "Distributed stochastic model predictive control synthesis for large-scale uncertain linear systems," in *Proc. Annual American Control Conf. (ACC)*, 2018, pp. 2071–2077.

[16] C. Conte, C. N. Jones, M. Morari, and M. N. Zeilinger, "Distributed synthesis and stability of cooperative distributed model predictive control for linear systems," *Automatica*, vol. 69, no. C, pp. 117–125, 2016.

[17] D. Q. Mayne, M. M. Seron, and S. V. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, vol. 41, no. 2, pp. 219 – 224, 2005.

[18] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1989.

[19] R. Rostami, G. Costantini, and D. Görges, "ADMM-based distributed model predictive control: Primal and dual approaches," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017, pp. 6598–6603.

[20] G. Calafiore and M. C. Campi, "Uncertain convex programs: randomized solutions and confidence levels," *Mathematical Programming*, vol. 102, no. 1, pp. 25–46, 2005.

[21] S. Muntwiler, K. P. Wabersich, A. Carron, and M. N. Zeilinger, "Distributed model predictive safety certification for learning-based control," *arXiv preprint arXiv:1911.01832*, 2019.

[22] R. B. Larsen, A. Carron, and M. N. Zeilinger, "Safe learning for distributed systems with bounded uncertainties," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 2536 – 2542, 2017, 20th IFAC World Congress.

[23] B. Recht, "A tour of reinforcement learning: The view from continuous control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, no. 1, pp. 253–279, 2019.

[24] F. Bullo, *Lectures on Network Systems*, 1st ed., 2019.