"Distributionally Robust Control of Constrained Stochastic Systems." Van Parys, Bart P.G., Daniel Kuhn, Paul J. Goulart, and Manfred Morari. IEEE Transactions on Automatic Control Vol. 61, No. 2 (2016): 430-442. DOI: 10.1109/TAC.2015.2444134.

Distributionally robust control of constrained stochastic systems

Bart P.G. Van Parys*, Daniel Kuhn†, Paul J. Goulart* and Manfred Morari*

Abstract

We investigate the control of constrained stochastic linear systems when faced with only limited information regarding the disturbance process, i.e. when only the first two moments of the disturbance distribution are known. We consider two types of *distributionally robust* constraints. The constraints of the first type are required to hold with a given probability for all disturbance distributions sharing the known moments. These constraints are commonly referred to as distributionally robust chance constraints with second-order moment specifications. In the second case, we impose conditional value-at-risk (CVaR) constraints to bound the expected constraint violation for all disturbance distributions consistent with the given moment information. Such constraints are referred to as distributionally robust CVaR constraints with second-order moment specifications. We argue that the design of linear controllers for systems with such constraints is both computationally tractable and practically meaningful for both finite and infinite horizon problems. The proposed methods are illustrated for a wind turbine blade control design case study where flexibility issues play an important role, and for which distributionally robust constraints constitute sensible design objectives.

I. Introduction

The problem of finding a control policy such that the state and inputs of an uncertain dynamical system remain in a given constraint set, despite the uncertain nature of the system, is an important and well studied problem within the control literature. In this article, we focus on discrete-time linear time-invariant (DLTI) systems, with system dynamics

$$\boldsymbol{x}_{t+1} = A\boldsymbol{x}_t + B\boldsymbol{u}_t + C\boldsymbol{w}_t,$$

*Automatic Control Laboratory, Swiss Federal Institute of Technology (ETH) Zürich, Physikstrasse 3, 8092 Zürich, CH

†Department of Computing, Imperial College London, 180 Queen's Gate, London SW7 2AZ, UK

where $\boldsymbol{x}_t$ is the state of the system, and $\boldsymbol{u}_t$ and $\boldsymbol{w}_t$ are the input and disturbance to the system, respectively. In constrained control problems, we are given a constraint set $\mathbb{X}$ for which we would like to express a requirement that "$\boldsymbol{x}_t \in \mathbb{X}$", i.e. the state $\boldsymbol{x}_t$ must remain in $\mathbb{X}$ *in some sense*. There are several common approaches to putting the design requirement "$\boldsymbol{x}_t \in \mathbb{X}$" on a mathematically sound footing. We first describe two standard methods for modeling such a constraint; the now classical worst-case approach and the more recent chance-constrained approach. Our notation is intentionally informal to start, with a more rigorous treatment deferred to later in the paper. We will propose two alternative approaches that overcome some of the deficiencies inherent in the two standard methods, particularly when faced with only a limited amount of information regarding the moments of the disturbance process $\boldsymbol{w}_t$.

Worst-case constraints: The problem of finding a control policy such that the state of an uncertain dynamical system remains in a given constraint set, for all possible disturbance realizations, is historically the most prevalent design goal within the control literature [3], [4], [23], [31]. This worst-case (or robust) formulation starts by assuming that the disturbance support is bounded and known, i.e. that the disturbance $\boldsymbol{w}_t$ is restricted to be an element of a bounded set W. The constraint "$\boldsymbol{x}_t \in \mathbb{X}$" is then interpreted as a condition that the state $\boldsymbol{x}_t$ is an element of $\mathbb{X}$ for all realizations of the disturbance process $\boldsymbol{w}_t$ generated from W. Identification of an optimal control policy for such problems is computationally intractable in general, so significant research effort has focussed on the development of design methods that provide admissible, but possibly suboptimal, control policies; see [16], [23] and the references therein.

The worst-case formulation requires that the support of the disturbance process is completely known and, in the presence of state constraints, a bounded set. This assumption may be quite restrictive, e.g. in cases where the disturbances are normally distributed and hence the disturbance support is unbounded. This motivates the need for approaches that make no such assumption regarding the support of the disturbance.

Chance constraints: Chance constraints require that the system's state constraints hold only with a specified probability level [11], [12]. The constraint "$\boldsymbol{x}_t \in \mathbb{X}$" is then modelled as

$$\mathbb{P}^\star \{\boldsymbol{x}_t \in \mathbb{X}\} \geq 1 - \epsilon, \tag{1}$$

where the probability measure $\mathbb{P}^\star$ is defined on the disturbance process $\boldsymbol{w}_t$ and is assumed known. Although no boundedness assumption is required on the disturbance support $\mathrm{supp}\{\mathbb{P}^\star\}$, in contrast to the worst-case approach, chance constraints are arguably worse from a practical perspective since they require the availability of a probability measure over the disturbances. Unfortunately, verifying a chance constraint in the form (1) is intractable under generic distributions, i.e. checking (1) for a given state distribution $\boldsymbol{x}_t$ is $\mathcal{NP}$-hard [25]. As a consequence, recent attention has shifted towards stochastic sampling methods, for which only probabilistic guarantees can typically be provided, e.g. that the chance constraint condition holds only with some level of confidence [8], [9].

In this paper we take an approach intermediate to these two extremes. Our goal is to provide a framework that addresses the constraint "$\boldsymbol{x}_t \in \mathbb{X}$" using only partial information about the *true but unknown* disturbance probability measure $\mathbb{P}^\star$. We briefly describe both of the new constraint models that will be introduced as alternatives to the worst-case and chance constrained problem formulations.

Distributionally robust chance constraints: In many situations the disturbance distribution $\mathbb{P}^\star$ is unknown and must be estimated from historical data, and hence is uncertain. We therefore assume only that the distribution $\mathbb{P}^\star$ belongs to a set $\mathcal{P}$ of distributions that share certain structural properties, i.e. their first two moments, which are assumed known. The distributionally robust counterpart [35] of the chance constraint (1) hence becomes

$$\forall \mathbb{P} \in \mathcal{P}: \quad \mathbb{P}\{\boldsymbol{x}_t \in \mathbb{X}\} \geq 1 - \epsilon. \tag{2}$$

The constraint (2) is referred to as a *distributionally robust chance constraint* [35] on $\boldsymbol{x}_t$ with a second moment specification. Such a constraint is a robust version of the classical chance constraint (1) in that it is insensitive to ambiguity in the disturbance distribution $\mathbb{P}^\star$, at least with respect to its higher order moments. One of the main advantages of this formulation over the classical chance constrained formulation is the fact that only partial information on the disturbance distribution $\mathbb{P}^\star$ is required.

Distributionally robust conditional value-at-risk (CVaR) constraints: For chance constraints in the form (2), a natural additional goal is to guarantee that the expected constraint violation in the remaining ϵ percent of the cases is small. To model such a requirement, we will consider distributionally robust CVaR constraints [34]. Our general approach is to measure, via some loss function L, the severity of constraint violation in the ϵ percent of the worst cases and to keep the degree of loss by this measure small. A precise statement of this distributionally robust CVaR approach requires some additional notation, and we defer the formalities to Section II.

We will show in the remainder of the paper how both the distributionally robust chance constraint and worst-case CVaR constraint interpretations of the condition "$\boldsymbol{x}_t \in \mathbb{X}$" constitute mathematically sound control design specifications. In particular, we will show that for optimal control problems with either constraint type, the resulting problems are both practically meaningful and computationally tractable. We stress that all of the numerical methods we present for solving such problems are deterministic, in contrast to the stochastic methods presented in [8], for which only probabilistic admissibility guarantees can be provided.

Outline: We provide a more mathematically rigorous description of distributionally robust chance and CVaR constraints in the context of control design problems in Section II, which can be read as an extended introduction. In Sections III and IV, we propose two control design problems, of finite and infinite horizon type respectively. We show that in addition to being practically justifiable, finding the optimal linear control policy in either case is a tractable problem when considered in conjunction with either of our alternative constraint descriptions. The latter of our proposed control problems is illustrated for a wind turbine blade control design case study in Section V, for which an assumption of limited moment information on the disturbance is quite natural. We have deliberately chosen a realistic and detailed numerical example in order to illustrate that the framework developed in the paper is not merely *l'art pour l'art*.

Notation and definitions

We denote by $\mathbb{I}_n$ the identity matrix in $\mathbb{R}^{n\times n}$ and by $\mathbb{S}^n_+$ and $\mathbb{S}^n_{++}$ the sets of all positive semidefinite and positive definite symmetric matrices in $\mathbb{R}^{n\times n}$, respectively. The diagonal concatenation of two matrices X and Y is denoted by $\mathrm{diag}(X,Y)$. All random vectors are defined as measurable functions on an abstract probability space $(\Omega, \mathcal{F}, \mathbb{P}^\star)$, where Ω is referred to as the sample space, $\mathcal{F}$ represents the σ-algebra of events, and $\mathbb{P}^\star$ denotes the true (but possibly unknown) probability measure. Without loss of generality we assume that Ω is sufficiently rich such that any (joint) distribution of all the random variables appearing in this paper on the Cartesian product of their individual range spaces is induced by a probability measure in $\mathcal{P}_0$.

This means that we can think of Ω as the Cartesian product of all the random variables' range spaces, in which case $\mathcal{F}$ is identified with the Borel σ-algebra on Ω, while each random variable reduces to a coordinate projection. For notational convenience, random vectors will be denoted in boldface, while their realizations will be denoted by the same symbols in normal font. For any $z \in \mathbb{R}$ we define $(z)^+ := \max\{z, 0\}$. The set $\mathcal{P}_0$ contains all probability measures on $(\Omega, \mathcal{F})$. For any random variable $\boldsymbol{x}$ we introduce the following shorthand notation

$$\mu_x := \mathbb{E}_{\mathbb{P}^\star}\left\{\boldsymbol{x}\right\}, \quad \Sigma_x := \mathbb{E}_{\mathbb{P}^\star}\left\{[\boldsymbol{x} - \mu_x] \cdot [\boldsymbol{x} - \mu_x]^\top\right\}, \quad M_x := \begin{pmatrix} \Sigma_x + \mu_x \mu_x^\top & \mu_x \\ \mu_x^\top & 1 \end{pmatrix}.$$

II. Distributionally robust CVaR and chance constraints

Chance constraints are popular as *soft constraints* that need only to hold up to a certain confidence level. Formally, the requirement that the n-dimensional random vector $\boldsymbol{x}$ should be contained in a set $\mathbb{X} \subseteq \mathbb{R}^n$ with high probability is expressed as

$$\mathbb{P}^\star(\boldsymbol{x} \in \mathbb{X}) \geq 1 - \epsilon, \tag{3}$$

where ϵ is a prescribed safety parameter that controls the level of acceptable constraint violations. Here, we could think of $\boldsymbol{x}$ as the random state of a linear dynamical system that depends both on previous control inputs (actions) and exogenous disturbances (noise). Chance constraints are often more practical than hard constraints, which can be viewed as degenerate chance constraints with $\epsilon = 0$ and which tend to encourage overly conservative decisions. Even worse, in linear dynamical systems hard state constraints typically become infeasible in the presence of unbounded (e.g. Gaussian) noise.

In spite of their conceptual appeal, chance constraints have not yet found wide application in optimization and control theory for a variety of reasons. On the one hand, the feasibility of a chance constraint can only be checked if the *true* distribution of the random vector $\boldsymbol{x}$ (which is determined by the true probability measure $\mathbb{P}^\star$) is precisely known. In practice, however, this distribution must almost invariably be estimated from noisy data and is therefore itself subject to ambiguity. This is problematic because even small changes in the distribution can have a dramatic impact on the geometry and size of the set of inputs or actions consistent with the chance constraint. Moreover, incorporating chance constraints into otherwise tractable optimization problems typically results in a non-convex problem, and consequently to computational intractability.

Finally, chance constraints bound the probability of constraint violation but do not impose any restrictions on the *degree* of the infeasibility encountered. However, severe constraint violations, i.e. scenarios in which the system state strays far outside of $\mathbb{X}$, are often much more harmful than mild violations in which the state remains close to the boundary of $\mathbb{X}$. Chance constraints fail to distinguish between these two situations and provide no mechanism to penalize severe constraint violations relative to mild ones.

In order to gain a better understanding of chance constraints, we require some terminology and notation. We will assume throughout that the set $\mathbb{X}$ is described as an intersection of lower level sets of finitely many convex loss functions $L_i : \mathbb{R}^n \to \mathbb{R}$, so that

$$\mathbb{X} := \{x \in \mathbb{R}^n \mid L_i(x) \leq 0 \ \forall i = 1, \ldots, I\}$$

and $\operatorname{int} \mathbb{X} = \{x \in \mathbb{R}^n \mid L_i(x) < 0 \ \forall i = 1, \ldots, I\}$. We refer to (3) as an *individual* chance constraint if $I = 1$ and as a *joint* chance constraint if $I > 1$. Every joint chance constraint can easily be reduced to an individual chance constraint by re-expressing $\mathbb{X}$ as $\{x \in \mathbb{R}^n \mid L(x; \alpha) \leq 0\}$, where the aggregate loss function

$$L(x; \alpha) := \max_{i=1,\ldots,I} \alpha_i L_i(x)$$

remains convex in x and depends on a set of strictly positive scaling parameters $\alpha \in \mathbb{R}^n_{++}$. Note that the particular choice of α has no impact on $\mathbb{X}$ and, consequently, no impact on the chance constraint (3). The reader may therefore regard α initially as a positive parameter that can be chosen arbitrarily. However, the flexibility to select α will be useful at a later stage to control the tightness of a tractable approximation of the chance constraint (3).

We now review an interesting connection between chance constraints and quantile-based risk measures that are commonly used in economics.

Definition II.1 (Value-at-risk)**.** *For any measurable loss function $L : \mathbb{R}^n \to \mathbb{R}$, probability measure $\mathbb{P}$ on $(\Omega, \mathcal{F})$ and tolerance $\epsilon \in (0, 1)$, the value-at-risk (VaR) of the random loss $L(\boldsymbol{x})$*

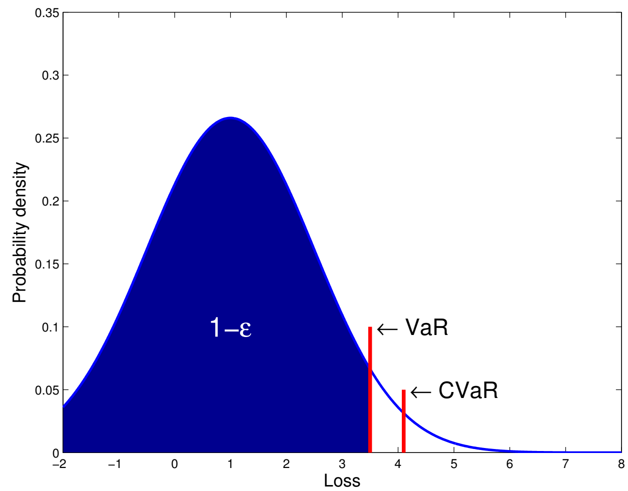

Fig. 1. Relation between VaR and CVaR for the illustrated loss distribution. The VaR is the $(1-\epsilon)$-quantile of the loss distribution, while the CVaR is the conditional expectation of loss above the $(1-\epsilon)$-quantile of the loss distribution, e.g. the CVaR is at the centre of mass of the ϵ-tail of the loss distribution.

at level ϵ with respect to $\mathbb{P}$ is defined as

$$\mathbb{P}\text{-VaR}_\epsilon\left(L(\boldsymbol{x})\right) := \inf\left\{\gamma \in \mathbb{R} \ \middle|\ \mathbb{P}(L(\boldsymbol{x}) > \gamma) \leq \epsilon\right\}.$$

We emphasize that the "value" at risk in this particular context is unrelated to the loss of economic currency, as in the usual interpretation in economics. In this constraint control context, "violation" at risk might be a more appropriate interpretation.

By definition, the VaR coincides with the $(1-\epsilon)$-quantile of the distribution of $L(\boldsymbol{x})$, as shown in Figure 1. Moreover, the reader may verify that the chance constraint (3) can be reformulated as a constraint on the VaR at level ϵ of the aggregate loss function $L(\boldsymbol{x};\alpha)$, that is,

$$\mathbb{P}^\star(\boldsymbol{x} \in \mathbb{X}) \geq 1-\epsilon \quad \iff \quad \mathbb{P}^\star(L(\boldsymbol{x};\alpha) > 0) \leq \epsilon \quad \iff \quad \mathbb{P}^\star\text{-VaR}_\epsilon\left(L(\boldsymbol{x};\alpha)\right) \leq 0. \tag{4}$$

Some useful properties of the VaR are summarized in the following lemma:

Lemma II.1 (Properties of VaR [14]). *Let $L : \mathbb{R}^n \to \mathbb{R}$ and $L' : \mathbb{R}^n \to \mathbb{R}$ be measurable loss functions, $\mathbb{P}$ a probability measure on $(\Omega, \mathcal{F})$ and $\epsilon \in (0,1)$. Then, the following hold:*

(i) *Monotonicity: $L(\boldsymbol{x}) \geq L'(\boldsymbol{x})$ $\mathbb{P}$-a.s. $\implies$ $\mathbb{P}\text{-VaR}_\epsilon(L(\boldsymbol{x})) \geq \mathbb{P}\text{-VaR}_\epsilon(L'(\boldsymbol{x}))$*

(ii) *Positive Homogeneity: $\lambda\mathbb{P}\text{-VaR}_\epsilon(L(\boldsymbol{x})) = \mathbb{P}\text{-VaR}_\epsilon(\lambda L(\boldsymbol{x}))$ for all $\lambda \in \mathbb{R}_+$*

(iii) *Translation Invariance:* $\mathbb{P}\text{-VaR}_\epsilon(L(\boldsymbol{x})+\lambda) = \mathbb{P}\text{-VaR}_\epsilon(L(\boldsymbol{x}))+\lambda$ *for all* $\lambda \in \mathbb{R}$.

A major deficiency of the VaR is its non-convexity in $L(\boldsymbol{x})$. In fact, it can be shown that $\mathbb{P}\text{-VaR}_\epsilon(L(\boldsymbol{x}))$ is generally non-convex in $\boldsymbol{x}$ even for linear loss functions. An alternative, convex, risk measure closely related to the VaR is the conditional value-at-risk defined next.

Definition II.2 (Conditional value-at-risk)**.** *For any measurable loss function* $L : \mathbb{R}^n \to \mathbb{R}$*, probability distribution* $\mathbb{P}$ *on* $(\Omega, \mathcal{F})$ *and tolerance* $\epsilon \in (0,1)$*, the CVaR of the random loss* $L(\boldsymbol{x})$ *at level* ϵ *with respect to* $\mathbb{P}$ *is defined as*

$$\mathbb{P}\text{-CVaR}_\epsilon\left(L(\boldsymbol{x})\right) := \inf_{\beta\in\mathbb{R}} \left\{ \beta + \frac{1}{\epsilon}\mathbb{E}_{\mathbb{P}}\left\{ (L(\boldsymbol{x})-\beta)^+ \right\} \right\}. \tag{5}$$

Rockafellar and Uryasev [29] have shown that the set of optimal solutions for β in (5) is a closed interval whose left endpoint is given by $\mathbb{P}\text{-VaR}_\epsilon\left(L(\boldsymbol{x})\right)$. Moreover, it can be shown that the CVaR admits the following equivalent representation [14, §4.4]:

$$\mathbb{P}\text{-CVaR}_\epsilon\left(L(\boldsymbol{x})\right) = \frac{1}{\epsilon}\int_0^\epsilon \mathbb{P}\text{-VaR}_\lambda\left(L(\boldsymbol{x})\right) \mathrm{d}\lambda. \tag{6}$$

This immediately implies that CVaR majorizes VaR, that is,

$$\mathbb{P}\text{-CVaR}_\epsilon\left(L(\boldsymbol{x})\right) \geq \mathbb{P}\text{-VaR}_\epsilon\left(L(\boldsymbol{x})\right). \tag{7}$$

If the random loss $L(\boldsymbol{x})$ follows a continuous distribution, then (6) can be integrated by parts to show that CVaR coincides with the conditional expectation of $L(\boldsymbol{x})$ above $\mathbb{P}\text{-VaR}_\epsilon\left(L(\boldsymbol{x})\right)$, as shown in Figure 1. This observation originally motivated the term *conditional value-at-risk*. The following lemma describes some basic properties of CVaR:

Lemma II.2 (Properties of CVaR [14])**.** *Let* $L : \mathbb{R}^n \to \mathbb{R}$ *and* $L' : \mathbb{R}^n \to \mathbb{R}$ *be measurable loss functions,* $\mathbb{P}$ *a probability measure on* $(\Omega, \mathcal{F})$ *and* $\epsilon \in (0,1)$*. Then, the following holds.*

(i) *Monotonicity:* $L(\boldsymbol{x}) \geq L'(\boldsymbol{x})$ $\mathbb{P}$*-a.s.* $\implies \mathbb{P}\text{-CVaR}_\epsilon(L(\boldsymbol{x})) \geq \mathbb{P}\text{-CVaR}_\epsilon(L'(\boldsymbol{x}))$

(ii) *Positive Homogeneity:* $\lambda\mathbb{P}\text{-CVaR}_\epsilon(L(\boldsymbol{x})) = \mathbb{P}\text{-CVaR}_\epsilon(\lambda L(\boldsymbol{x}))$ *for all* $\lambda \in \mathbb{R}_+$

(iii) *Translation Invariance:* $\mathbb{P}\text{-CVaR}_\epsilon(L(\boldsymbol{x})+\lambda) = \mathbb{P}\text{-CVaR}_\epsilon(L(\boldsymbol{x}))+\lambda$ *for all* $\lambda \in \mathbb{R}$

(iv) *Convexity: For each* $\lambda \in [0,1]$ *we have*

$$\lambda\mathbb{P}\text{-CVaR}_\epsilon\left(L(\boldsymbol{x})\right) + (1-\lambda)\mathbb{P}\text{-CVaR}_\epsilon\left(L'(\boldsymbol{x})\right) \geq \mathbb{P}\text{-CVaR}_\epsilon\left(\lambda L(\boldsymbol{x}) + (1-\lambda)L'(\boldsymbol{x})\right).$$

One of the main reasons for the popularity of CVaR is the fact that it represents a convex and conservative (i.e. pessimistic) approximation of VaR. Recalling that chance constraints can always be rewritten as VaR constraints, this observation prompts us to use CVaR constraints as conservative approximations for chance constraints. Indeed, (4) and (7) imply

$$\mathbb{P}^\star\text{-CVaR}_\epsilon\left(L(\boldsymbol{x};\alpha)\right) \leq 0 \quad \Longrightarrow \quad \mathbb{P}^\star\text{-VaR}_\epsilon\left(L(\boldsymbol{x};\alpha)\right) \leq 0 \quad \Longleftrightarrow \quad \mathbb{P}^\star(\boldsymbol{x} \in \mathbb{X}) \geq 1-\epsilon. \tag{8}$$

Note that for convex loss functions the set of all random vectors $\boldsymbol{x}$ satisfying the CVaR constraint in (8) is convex due to the convexity and monotonicity of CVaR.

In economic theory, CVaR traditionally measures an economic loss, hence the function L is given ab initio. In control practice however, one is typically given a constraint set $\mathbb{X}$, and is free to select any loss functions L_i compatible with $\mathbb{X}$, i.e. one can choose any L_i satisfying $\mathbb{X} = \{x \in \mathbb{R}^n \mid L(x;\alpha) \leq 0\}$. The choice of weights α_i can then be used to indicate the relative importance of the individual loss functions L_i.

CVaR constraints address two of the main shortcomings of chance constraints. First, unlike chance constraints, they lead to tractable convex optimization problems. Second, CVaR constraints impose a higher penalty on realizations of $\boldsymbol{x}$ that materialize far outside of $\mathbb{X}$ (i.e. with $L(\boldsymbol{x};\alpha) \gg 0$) and therefore penalize severe constraint violations more aggressively than mild ones. In contrast, chance constraints impose uniform penalties on all constraint violations irrespective of their degree of infeasibility.

Unfortunately, checking the feasibility of CVaR constraints still requires precise knowledge of the *true* probability measure $\mathbb{P}^\star$. In practice, only limited information about $\mathbb{P}^\star$ may be available, such as the support or some descriptive measures of the location and dispersion of certain random variables under $\mathbb{P}^\star$. Abstractly, we can represent the limited available information about $\mathbb{P}^\star$ by a set $\mathcal{P}$ of probability measures on $(\Omega, \mathcal{F})$ with the following properties: (i) It is known that $\mathbb{P}^\star \in \mathcal{P}$, and (ii) $\mathcal{P}$ is the smallest set of probability distributions for which we can guarantee that $\mathbb{P}^\star \in \mathcal{P}$. We will henceforth refer to $\mathcal{P}$ as an *ambiguity set*.

To immunize the chance constraint (3) against distributional ambiguity, we may require that it should hold for each probability measure in the ambiguity set. The resulting *distributionally robust chance constraint* can be represented as

$$\mathbb{P}(\boldsymbol{x} \in \mathbb{X}) \geq 1-\epsilon \quad \forall \mathbb{P} \in \mathcal{P} \quad \Longleftrightarrow \quad \inf_{\mathbb{P}\in\mathcal{P}} \mathbb{P}(\boldsymbol{x} \in \mathbb{X}) \geq 1-\epsilon. \tag{9}$$

Similarly, recalling that $\mathbb{X} = \{x \in \mathbb{R}^n \mid L(x;\alpha) \leq 0\}$ for any $\alpha \in \mathbb{R}_{++}$, we can immunize the CVaR constraint on the left hand side of (8) against distributional ambiguity. The resulting *distributionally robust CVaR constraint* takes the form

$$\mathbb{P}\text{-CVaR}_\epsilon(L(\boldsymbol{x};\alpha)) \leq 0 \ \forall \mathbb{P} \in \mathcal{P} \iff \sup_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon(L(\boldsymbol{x};\alpha)) \leq 0. \tag{10}$$

As in the classical setting without distributional ambiguity, it can be shown that (10) provides a conservative approximation for (9). Indeed, it is easy to see that the worst-case CVaR dominates the worst-case VaR, and thus

$$\sup_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon(L(\boldsymbol{x};\alpha)) \leq 0 \implies \sup_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-VaR}_\epsilon(L(\boldsymbol{x};\alpha)) \leq 0 \iff \inf_{\mathbb{P}\in\mathcal{P}} \mathbb{P}(\boldsymbol{x} \in \mathbb{X}) \geq 1-\epsilon,$$

see, e.g. [13] or [35].

In order to further characterize the relation between worst-case VaR and worst-case CVaR, and to facilitate statements about computational tractability, we require some structural assumptions about the ambiguity set $\mathcal{P}$ and the loss functions $L_i(x)$, $i = 1, \ldots, I$ defining the constraint set $\mathbb{X}$. We henceforth assume that the ambiguity set $\mathcal{P}$ contains all probability measures $\mathbb{P}$ under which the random variable $\boldsymbol{x}$ has a given mean value $\mu_x \in \mathbb{R}^n$ and a given covariance matrix $\Sigma_x \in \mathbb{S}^n_{++}$. Moreover, we require that each constraint function L_i is convex and quadratic, and is representable as $L_i(x) = x^\top E_i x + 2e_i^\top x + e_i^0$ for some $E_i \in \mathbb{S}^n_+$, $e_i \in \mathbb{R}^n$ and $e_i^0 \in \mathbb{R}$. This means that $\mathbb{X}$ can be any intersection of half-spaces and generalized ellipsoids. The aggregate loss $L(x;\alpha)$ is then given by a pointwise maximum of quadratic functions.

We next recall some tractability and exactness results relating to the CVaR approximation.

Theorem II.1 (Tractability of worst-case CVaR [35]). *Under the preceding assumptions about the ambiguity set $\mathcal{P}$ and the loss functions L_i, $i = 1, \ldots, I$, the worst-case CVaR is equivalent to the following tractable semi-definite program (SDP):*

$$\begin{aligned}
\sup_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon(L(\boldsymbol{x};\alpha)) \ = \ \inf_{\beta,X} \ & \beta + \frac{1}{\epsilon}\mathrm{Tr}\left\{M_x X\right\} \\
\text{s.t.} \ & X \in \mathbb{S}^{n+1}_+, \ \beta \in \mathbb{R} \\
& X - \begin{pmatrix} \alpha_i E_i & \alpha_i e_i \\ \alpha_i e_i^\top & \alpha_i e_i^0 - \beta \end{pmatrix} \succeq 0, \quad \forall i \in \{1, \ldots, I\}.
\end{aligned} \tag{11}$$

If there is only one loss function that is concentric with the distribution of $\boldsymbol{x}$, then the SDP (11) admits an analytical solution. So far, this special case has not been considered in the literature. However, we will see in Section IV that it is relevant for constrained infinite horizon problems.

Corollary II.1 (Concentric distributions and loss functions)**.** *If $\mathbb{X}$ constitutes a single ellipsoid centred at the origin (i.e. $I = 1$ and $e_1 = 0$), while the random vector $\boldsymbol{x}$ has mean $\mu_x = 0$, then*

$$\sup_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon\left(L_1(\boldsymbol{x})\right) = e_1^0 + \frac{1}{\epsilon}\mathrm{Tr}\left\{\Sigma_x E_1\right\}. \tag{12}$$

Proof: For $I = 1$, $\alpha_1 = 1$ and $e_1 = \mu_x = 0$ Theorem II.1 implies

$$\begin{aligned}
\sup_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon(L_1(\boldsymbol{x})) \;=\; & \inf \;\; \beta + \frac{1}{\epsilon}\left(\mathrm{Tr}\left\{\Sigma_x Y\right\} + y_0\right) \\
& \text{s.t.} \;\; Y \in \mathbb{S}^n_+,\; y \in \mathbb{R}^n,\; y_0 \in \mathbb{R}_+,\; \beta \in \mathbb{R} \\
& \qquad \begin{pmatrix} Y & y \\ y^\top & y_0 \end{pmatrix} \succeq 0, \quad \begin{pmatrix} Y - E_1 & y \\ y^\top & y_0 - e_1^0 + \beta \end{pmatrix} \succeq 0.
\end{aligned} \tag{13}$$

As $Y = E_1$, $y = 0$, $y_0 = 0$ and $\beta = e_1^0$ is feasible in (13), it is clear that the worst-case CVaR is bounded above by $e_1^0 + \frac{1}{\epsilon}\mathrm{Tr}\left\{\Sigma_x E_1\right\}$. To prove the converse inequality, we let $(Y^\star, y^\star, y_0^\star, \beta^\star)$ be an optimal solution of (13). Then, we find

$$\begin{aligned}
\sup_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon\left(L_1(\boldsymbol{x})\right) \;&=\; \beta^\star + \frac{1}{\epsilon}\left(\mathrm{Tr}\left\{\Sigma_x Y^\star\right\} + y_0^\star\right) \\
&\geq\; \beta^\star + \frac{1}{\epsilon}\Big(\mathrm{Tr}\left\{\Sigma_x E_1\right\} + (e_1^0 - \beta^\star)^+\Big) \;\geq\; e_1^0 + \frac{1}{\epsilon}\mathrm{Tr}\left\{\Sigma_x E_1\right\},
\end{aligned}$$

where the first inequality exploits the feasibility of $(Y^\star, y^\star, y_0^\star, \beta^\star)$ in (13), and the second inequality exploits the fact that $y_0^\star \geq (e_1^0 - \beta^\star)^+$ and $\epsilon \in (0, 1)$. ∎

Theorem II.1 implies that the CVaR constraint $\sup_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon(L(\boldsymbol{x};\alpha)) \leq 0$ is equivalent to a collection of linear matrix inequality (LMI) constraints. Similarly, under the assumptions of Corollary II.1 the CVaR constraint $\sup_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon(L_1(\boldsymbol{x})) \leq 0$ is equivalent to a single scalar linear inequality.

Even though these CVaR constraints generally provide *conservative* approximations for the corresponding distributionally robust chance constraints, the approximations become *essentially exact*[1] for a judicious choice of the scaling parameters $\alpha \in \mathbb{R}^n_{++}$.

[1]Up to an interior operation, see equivalence (14).

Theorem II.2 (Exactness of the CVaR Approximation [7], [35]). *Under the above assumptions about the ambiguity set $\mathcal{P}$ and the loss functions L_i, $i = 1, \ldots, I$, the worst-case CVaR constraint is equivalent to a variant of a distributionally robust chance constraint where $\mathbb{X}$ is replaced with its interior if we can optimize over $\alpha \in \mathbb{R}^n_{++}$, i.e.*

$$\inf_{\alpha \in \mathbb{R}^n_{++}} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon(L(\boldsymbol{x}; \alpha)) \leq 0 \iff \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\boldsymbol{x} \in \operatorname{int} \mathbb{X}) \geq 1 - \epsilon. \tag{14}$$

If the set $\mathbb{X}$ is described by a single quadratic inequality (i.e. $I = 1$), then the worst-case CVaR constraint involving L_1 (instead of the aggregate loss function) is equivalent to the distributionally robust chance constraint, i.e.

$$\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon(L_1(\boldsymbol{x})) \leq 0 \iff \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\boldsymbol{x} \in \mathbb{X}) \geq 1 - \epsilon. \tag{15}$$

Note that the inclusion $\operatorname{int} \mathbb{X} \subseteq \mathbb{X}$ implies

$$\inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\boldsymbol{x} \in \operatorname{int} \mathbb{X}) \leq \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\boldsymbol{x} \in \mathbb{X}),$$

but except for degenerate situations the two worst-case probabilities in the above expression are equal. Theorems II.1 and II.2 thus imply that the worst-case chance constraint is essentially equivalent to the following constraints:

$$\exists X \in \mathbb{S}^{n+1}_+,\ \beta \in \mathbb{R},\ \alpha \in \mathbb{R}^I_{++} :\ \beta + \frac{1}{\epsilon} \operatorname{Tr}\{M_x X\} \leq 0,\ X - \begin{pmatrix} \alpha_i E_i & \alpha_i e_i \\ \alpha_i e_i^\top & \alpha_i e_i^0 - \beta \end{pmatrix} \succeq 0\ \forall i \in \{1, \ldots, I\}.$$

III. Finite horizon distributionally robust control problems

We consider a DLTI system with n states, m control inputs, r outputs, d exogenous inputs or disturbances and r measurements:

$$\begin{cases} \boldsymbol{x}_{t+1} = A\boldsymbol{x}_t + B\boldsymbol{u}_t + C\boldsymbol{w}_t \quad \text{and} \quad \boldsymbol{x}_0 = x \\ \boldsymbol{y}_t = D\boldsymbol{x}_t + E\boldsymbol{w}_t, \end{cases} \tag{$\mathcal{S}$}$$

where all matrices are of appropriate dimension and the disturbances $\boldsymbol{w}_t$ model both process noise (via the term $C\boldsymbol{w}_t$) and measurement noise (via $E\boldsymbol{w}_t$). The input $\boldsymbol{u}_t$ is restricted to be $\mathcal{F}^y_t := \sigma(\boldsymbol{y}_0, \ldots, \boldsymbol{y}_t)$ -measurable. Our goal is to design a finite-horizon control law for the system ($\mathcal{S}$) that minimizes an expected value quadratic cost, subject to an additional requirement that the state satisfies the constraint "$\boldsymbol{x}_t \in \mathbb{X}$" according to either a chance-constrained or CVaR interpretation. We wish to do this despite some ambiguity on the disturbance distribution.

Specifically, we assume only that the following information is available about the disturbance process:

Assumption III.1 (Weak sense stationary disturbances)**.** *We assume that in the DLTI system $\mathcal{S}$, the disturbance $\boldsymbol{w}_t$ is a weak sense stationary (w.s.s.) white noise process with covariance matrix $\Sigma_{w_t} = \mathbb{I}_d$[2] and mean $\mu_{w_t} = \mu$ for all $t \in \mathbb{N}_0$.*

The w.s.s. assumption appears frequently in signal processing [28], but is less common in the control literature. In effect, it assumes that only the autocorrelation $R_{ww}(t) := \mathbb{E}_{\mathbb{P}^\star}\{\boldsymbol{w}_i \cdot \boldsymbol{w}_{i-t}^\top\}$ is known, with $R_{ww}(0) = \mathbb{I}_d + \mu\mu^\top$ and $R_{ww}(t) = \mu\mu^\top$ otherwise. Furthermore, knowing the first two moments of a w.s.s. process is, by merit of the Wiener-Khintchine Theorem, equivalent to knowing its power spectrum [28]. Estimating the spectral density of the disturbance $\boldsymbol{w}_t$, for example from historical data[3], is significantly easier in practice than determining the complete distribution $\mathbb{P}^\star$.

The w.s.s. assumption implies that the only information available about the disturbance distribution $\mathbb{P}^\star$ is its autocorrelation function. Hence, $\mathbb{P}^\star$ is only known to be an element of the ambiguity set

$$\mathcal{P}_\infty := \left\{ \mathbb{P} \in \mathcal{P}_0 \;\middle|\; \mathbb{E}_{\mathbb{P}}\left\{(\boldsymbol{w}_i^\top, 1)^\top \cdot (\boldsymbol{w}_j^\top, 1)^\top\right\} = \begin{pmatrix} \mathbb{I}_d\delta_{ij} + \mu\mu^\top & \mu \\ \mu^\top & 1 \end{pmatrix}, \quad \forall i, j \in \mathbb{N}_0 \right\}.$$

The set $\mathcal{P}_\infty$ contains all probability distributions consistent with the known moment information. When choosing a control policy for the system $\mathcal{S}$, we will require that it be distributionally robust with respect to the ambiguity set $\mathcal{P}_\infty$, in either a chance constrained or CVaR sense, for the constraint "$\boldsymbol{x}_t \in \mathbb{X}$". In order to achieve this control design objective, the notion of a distributionally robust constraint, introduced in Section II, is now used to formulate our control problem.

Control constraints: We will consider distributionally robust constraints for the system $\mathcal{S}$

[2] The assumption that the covariance of the disturbance $\boldsymbol{w}_t$ is the identity matrix $\mathbb{I}_d$, is without loss of generality. In case Σ_{w_t} is not the identity, taking the Cholesky decomposition $\Sigma_{w_t} = LL^\top$ and substituting $C \leftarrow CL$, $E \leftarrow EL$ we can obtain an equivalent system which satisfies Assumption III.1.

[3] The study of this problem is referred to as *spectral density estimation* in the signal processing community.

enforced over a finite time horizon $t \in \{0, \dots, N\}$, i.e.

$$\sup_{\mathbb{P} \in \mathcal{P}_\infty} \mathbb{P}\text{-}\mathbf{CVaR}_\epsilon \left(L \left(\boldsymbol{x}_t; \alpha \right) \right) \leq 0 \quad \forall t \in \{0, \dots, N-1\}. \tag{16}$$

We will refer to the parameter N as the horizon length of the finite optimal control problem studied. We assume that the parameter $\alpha \in \mathbb{R}^n_{++}$ is given, either as an attempt to approximate a distributionally robust chance constraint or as an indicator of the relative importance of the loss severity measures L_i.

Assumption III.2. *An aggregated loss function $L : \mathbb{R}^n \times \mathbb{R}^n_{++} \to \mathbb{R}$ for the distributionally robust CVaR constraints* (16) *is given as $L(x; \alpha) = \max_{i \in \{1, \dots, I\}} \left[\alpha_i \left(x^\top E_i x + 2 e_i^\top x + e_i^0 \right) \right]$ where $E_i \in \mathbb{S}^n_+$, $e_i \in \mathbb{R}_n$, $e_i^0 \in \mathbb{R}$.*

Hence, the set $\mathbb{X}$ corresponding to a loss function satisfying Assumption III.2 is a finite intersection of half-spaces and generalized ellipsoids.

For the system $\mathcal{S}$ we define a causal control policy $\pi_N := \{\boldsymbol{u}_0, \boldsymbol{u}_1, \dots, \boldsymbol{u}_{N-1}\}$, such that the control input selected at each time $t \in [0, \dots, N-1]$ is a function mapping prior measurements to actions, i.e. $\boldsymbol{u}_t$ is $\mathcal{F}_t^y$-measurable, where we assume that the initial state $\boldsymbol{x}_0 = x$ is known without any loss of generality[4]. We denote the set of all such policies as Π_N. We wish to find, if it exists, a policy $\pi_N \in \Pi_N$ such that system $\mathcal{S}$ satisfies the CVaR constraints (16) over a finite horizon. We refer to such a policy as *admissible* with respect to the system $\mathcal{S}$ and the CVaR constraints (16).

Objective function: Our aim is to find a causal control policy $\pi_N \in \Pi_N$ that is admissible with respect to the CVaR constraints while minimizing a given objective function J_N. We will assume throughout that the objective function $J_N : \mathbb{R}^n \times \Pi_N \to \mathbb{R}_+$ is a discounted sum of quadratic stage costs, i.e. that it is in the form

$$J_N(x, \pi_N) := \sup_{\mathbb{P} \in \mathcal{P}_\infty} \mathbb{E}_\mathbb{P} \left\{ \sum_{t=0}^{N-1} \beta^t \left[\boldsymbol{x}_t^\top Q \boldsymbol{x}_t + \boldsymbol{u}_t^\top R \boldsymbol{u}_t \right] + \beta^N \boldsymbol{x}_N^\top Q_f \boldsymbol{x}_N \right\}, \tag{17}$$

where we refer to $\beta \in (0, 1]$ as the discount factor of the control cost. It is assumed that the objective function J_N is convex, i.e. Q, $Q_f \in \mathbb{S}_+$ and $R \in \mathbb{S}_{++}$. We are therefore interested in

[4]In the case that the initial state $\boldsymbol{x}_0 = \boldsymbol{x}$ is itself uncertain, one can always add an additional leading state $\boldsymbol{x}_{-1} = 0$ and a state update equation $\boldsymbol{x}_0 = A\boldsymbol{x}_{-1} + \boldsymbol{x}$, where $\boldsymbol{x}$ would be a noise term.

the solution to the optimal control problem

$$\begin{aligned}\inf_{\pi_N \in \Pi_N} \quad & J_N(x, \pi_N) \\ \text{s.t.} \quad & \boldsymbol{x}_{t+1} = A\boldsymbol{x}_t + B\boldsymbol{u}_t + C\boldsymbol{w}_t,\ \boldsymbol{x}_0 = x \\ & \sup_{\mathbb{P}\in\mathcal{P}_\infty} \mathbb{P}\text{-}\mathbf{CVaR}_\epsilon\left(L\left(\boldsymbol{x}_t;\alpha\right)\right) \leq 0, \quad \forall t \in \{0,\ldots,N-1\}\end{aligned} \tag{$\mathcal{R}_N$}$$

This problem appears to be intractable in general, since; (i) optimizing directly over arbitrary measurable policies π_N in Π_N seems to be out of the question; and (ii) distributionally robust constraints such as (16), even for convex loss functions L, seem hard to deal with directly when $\boldsymbol{x}_t$ is a general non-linear function.

Hence, in what follows we restrict attention to control policies that are affine in the past disturbances as in [15]. Restricted policies of this type are well known in the operations research and control community, where they are commonly referred to either as *linear decision rules* [2] or *affine feedback policies* [16]. Although such policies are typical suboptimal, recent research effort has focussed on providing suboptimality bounds when applied to systems with worst-case constraints [17], [18], [27].

Denote by $\boldsymbol{x} := (\boldsymbol{x}_0^\top,\ldots,\boldsymbol{x}_N^\top)^\top$, $\boldsymbol{u} := (\boldsymbol{u}_0^\top,\ldots,\boldsymbol{u}_{N-1}^\top)^\top$ and $\boldsymbol{y} := (\boldsymbol{y}_0^\top,\ldots,\boldsymbol{y}_{N-1}^\top)^\top$ the collection of states, inputs and measurements, respectively, over the given finite horizon. We similarly define a vector of disturbances as

$$\boldsymbol{w} := (1, \boldsymbol{w}_0^\top, \ldots, \boldsymbol{w}_{N-1}^\top)^\top, \tag{18}$$

augmented with a leading one. This leading term is included for notational convenience so that any affine function of $(\boldsymbol{w}_0,\ldots,\boldsymbol{w}_{N-1})$ can be written as $X\boldsymbol{w}$ for some matrix X with appropriate dimensions. Because of the w.s.s. condition on the disturbance process in Assumption III.1, we have that $\mathbb{E}_{\mathbb{P}^\star}\left\{\boldsymbol{w}\cdot\boldsymbol{w}^\top\right\} = M_w \in \mathbb{S}_{++}^{Nd+1}$ with

$$M_w := \left(1, \mu^\top, \ldots, \mu^\top\right)^\top \left(1, \mu^\top, \ldots, \mu^\top\right) + \operatorname{diag}\left(0, \mathbb{I}_N \otimes \mathbb{I}_d\right).$$

The dynamics of the linear system $\mathcal{S}$ over the finite horizon N can then be written as

$$\mathbf{x} = \mathcal{B}\mathbf{u} + \mathcal{C}\mathbf{w},\ \boldsymbol{y} = \mathcal{D}\boldsymbol{u} + \mathcal{E}\boldsymbol{w}, \tag{19}$$

for some matrices $(\mathcal{B},\mathcal{C},\mathcal{D},\mathcal{E})$ that can be derived from the system matrices and initial state $\boldsymbol{x}_0 = x$; see Appendix A. Note in particular that the leading one in (18) means that the term $\mathcal{C}\mathbf{w}$

is an affine function of both the disturbances and the initial state $\boldsymbol{x}_0 = x$. Our approach will be to restrict $\boldsymbol{u}$ to be affine in the past disturbances, i.e. $\boldsymbol{u} = U\boldsymbol{w}$ for some causal feedback matrix $U \in \mathcal{N}$.

The causality set $\mathcal{N}$ must ensure that the resulting feedback policy $\boldsymbol{u}_t$ is $\mathcal{F}_t^y$-measurable, i.e. that the feedback policy $\boldsymbol{u}_t$ depends only on the initial state x and observed outputs $[\boldsymbol{y}_0, \dots, \boldsymbol{y}_t]$. This can be achieved by a reparametrization of the feedback policy $\boldsymbol{u} = \tilde{U}\boldsymbol{\eta}$ as an affine function of the *purified observations* $\boldsymbol{\eta} = (\mathcal{DC} + \mathcal{E})\,\boldsymbol{w}$ as discussed in [2, §14.4.2]. The causality set can then be defined as

$$\mathcal{N} := \left\{ U \in \mathbb{R}^{N_x \times N_w} \;\middle|\; U = \begin{pmatrix} u_0 & 0 & 0 & 0 & 0 & 0 \\ u_1 & U_{1,0} & 0 & 0 & 0 & 0 \\ u_2 & U_{2,0} & U_{2,1} & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & \vdots \\ u_{N-2} & U_{N-2,0} & U_{N-2,1} & \dots & U_{N-2,N-2} & 0 \\ u_{N-1} & U_{N-1,0} & U_{N-1,1} & \dots & U_{N-1,N-2} & U_{N-1,N-1} \end{pmatrix} (\mathcal{DC} + \mathcal{E}) \right\}$$

which ensures that $\boldsymbol{u}_t$ is $\mathcal{F}_t^y$-measurable. Assume we have such an affine policy $\boldsymbol{u} = U\boldsymbol{w}$, then the cost of this policy according to the cost function (17) is

$$\tilde{J}_N(x, U) := \mathrm{Tr}\left\{ U^\top (J_u + \mathcal{B} J_x \mathcal{B})\, U M_w + 2\mathcal{C} J_x \mathcal{B} U M_w + \mathcal{C}^\top J_x \mathcal{C} M_w \right\},$$

where $J_x := \mathrm{diag}\big(\mathrm{diag}\left(\beta^0, \dots, \beta^{N-1}\right) \otimes Q, \beta^N Q_f\big)$ and $J_u := \mathrm{diag}\left(\beta^0, \dots, \beta^{N-1}\right) \otimes R$. Note that $\tilde{J}_N(x, U)$ is convex quadratic in U since $\mathrm{diag}(Q, R) \in \mathbb{S}_+$. We are now ready to state the main result of this section, which shows that finding the best affine control policy for problem $\mathcal{R}_N$ can be reformulated as a tractable convex optimization problem.

Theorem III.1 (CVaR constrained control)**.** *The best admissible affine control policy of problem $\mathcal{R}_N$, i.e. a solution to the restricted problem*

$$\begin{aligned} \inf_{U \in \mathcal{N}} \quad & \tilde{J}_N(x, U) \\ \text{s.t.} \quad & \mathbf{x} = \mathcal{B}\mathbf{u} + \mathcal{C}\mathbf{w}, \;\; \boldsymbol{u} = U\boldsymbol{w} \\ & \sup_{\mathbb{P} \in \mathcal{P}_\infty} \mathbb{P}\text{-CVaR}_\epsilon\left(L(\boldsymbol{x}_t; \alpha)\right) \le 0, \quad \forall t \in \{0, \dots, N-1\} \end{aligned} \qquad (\tilde{\mathcal{R}}_N)$$

where the loss function $L : \mathbb{R}^n \times \mathbb{R}^I_{++} \to \mathbb{R}$ satisfies Assumption III.2, can be found as a solution *of the SDP*

$$
\begin{array}{ll}
\inf & \tilde{J}_N(x,U) \\
\text{s.t.} & \left.\begin{array}{l}
U \in \mathcal{N},\ \beta_t \in \mathbb{R},\ X_t \in \mathbb{S}_+^{Nd+2},\ P_t^i \in \mathbb{S}_+^{Nd+1} \\
\beta_t + \dfrac{1}{\epsilon}\mathrm{Tr}\left\{M_w X_t\right\} \le 0, \\
X_t - \begin{pmatrix} \alpha_i P_t^i & \alpha_i(\mathcal{B}_t U + \mathcal{C}_t)^\top e_i \\ e_i^\top(\mathcal{B}_t U + \mathcal{C}_t)\alpha_i & \alpha_i e_i^0 - \beta_t \end{pmatrix} \succeq 0, \\
\begin{pmatrix} P_t^i & (\mathcal{B}_t U + \mathcal{C}_t)^\top E_i^{1/2} \\ E_i^{1/2}(\mathcal{B}_t U + \mathcal{C}_t) & \mathbb{I}_n \end{pmatrix} \succeq 0,
\end{array}\right\} \begin{array}{l} \forall t \in \{0,\dots,N-1\} \\ \forall i \in \{1,\dots,I\} \end{array}
\end{array} \tag{20}
$$

where $\mathcal{B} =: (\mathcal{B}_0^\top, \dots, \mathcal{B}_{N-1}^\top)^\top$ *and* $\mathcal{C} =: (\mathcal{C}_0^\top, \dots, \mathcal{C}_{N-1}^\top)^\top$.

Proof: See Appendix A. ■

We remark that the equivalence (14) in Theorem II.2 ensures that there exists some $\alpha \in \mathbb{R}_{++}$ such that the constraints (16) reduce to a distributionally robust chance constraint for $\mathbb{X} = \{x \mid L(x;\alpha) \le 0\}$, whenever $\boldsymbol{x}_t$ is an affine function of the disturbances. In principle one could therefore identity such a parameter vector α to recover an exact representation of a robust chance constraint in the problem ($\tilde{\mathcal{R}}_N$). However, simultaneous optimization over both U and α in (20) would result in a non-convex bi-affine optimization problem, and such problems are known to be $\mathcal{NP}$-hard in general.

While Theorem II.2 is only existential in nature, i.e. it is true for some unknown $\alpha \in \mathbb{R}_{++}$, the equivalence between chance constraints and CVaR constraints when $\mathbb{X}$ is a simple ellipsoid or $I = 1$ is guaranteed. This result enables us to formulate the following corollary to Theorem III.1.

Corollary III.1 (Chance constrained control)**.** *The best admissible affine control policy of the restricted problem*

$$
\begin{array}{ll}
\inf\limits_{U \in \mathcal{N}} & \tilde{J}_N(x,U) \\
\text{s.t.} & \mathbf{x} = \mathcal{B}\mathbf{u} + \mathcal{C}\mathbf{w},\ \boldsymbol{u} = U\boldsymbol{w} \\
& \inf\limits_{\mathbb{P} \in \mathcal{P}_\infty} \mathbb{P}\left\{\boldsymbol{x}_t \in \mathbb{X}\right\} \ge 1 - \epsilon,\ \forall t \in \{0,\dots,N-1\}
\end{array}
$$

where the constraint set $\mathbb{X} = \left\{x \;\middle|\; x^\top E_1 x + 2e_1^\top x + e_1^0 \leq 0\right\}$ *is a single ellipsoid, can be found as a solution of the SDP* (20) *with* $I = 1$ *and* $\alpha_1 = 1$.

IV. Infinite horizon distributionally robust control problems

Infinite horizon control problems lend themselves to applications in which transient behaviour is of lesser importance, but in which we are interested in steady state behaviour. In Section V we present a numerical example of such a problem in the context of wind turbine blade control.

The problem setting is similar to the one presented in Section III, in that we again consider the DLTI system $\mathcal{S}$ where the disturbance input process $\boldsymbol{w}_t$ satisfies Assumption III.1. In addition, we assume that the disturbance $\boldsymbol{w}_t$ has zero mean $\mu_{w_t} = \mu = 0$, a zero initial condition $\boldsymbol{x}_0 = 0$ reflects our indifference towards transient behaviour.

Control constraints : We will consider the problem of finding a *causal* linear time-invariant feedback law π, that satisfies the following limit or steady state constraint,

$$\limsup_{t\to\infty} \quad \sup_{\mathbb{P}\in\mathcal{P}_\infty} \mathbb{P}\text{-}\mathbf{CVaR}_\epsilon\left(L_0\left(\boldsymbol{x}_t\right)\right) \leq 0, \tag{21}$$

where the loss function $L_0(x) := x^\top E_0 x + e^0$ is a convex quadratic function with $E_0 \in \mathbb{S}^n_{++}$. This implies that the corresponding constraint set $\mathbb{X} := \{x \mid L_0(x) \leq 0\}$ is an ellipsoid centred at the origin. The set of all linear time-invariant and causal feedback policies satisfying constraint (21) is denoted by Π_∞.

Objective function: We are interested in finding a feedback law $\pi \in \Pi_\infty$ that minimizes the infinite horizon limit of the stage cost function in (17) for system $\mathcal{S}$, with no discounting or terminal cost. The design goal in this case reduces to minimizing the average stage cost, so that the objective function becomes

$$\begin{aligned} J_\infty(\pi) &:= \limsup_{N\to\infty} \quad \sup_{\mathbb{P}\in\mathcal{P}_\infty} \frac{1}{N}\mathbb{E}_{\mathbb{P}}\left\{\sum_{t=0}^{N-1}\left[\boldsymbol{x}_t^\top Q\boldsymbol{x}_t + \boldsymbol{u}_t^\top R\boldsymbol{u}_t\right]\right\}, \\ &= \limsup_{N\to\infty} \quad \frac{1}{N}\sum_{t=0}^{N-1}\mathbb{E}_{\mathbb{P}^\star}\left\{\boldsymbol{x}_t^\top Q\boldsymbol{x}_t + \boldsymbol{u}_t^\top R\boldsymbol{u}_t\right\}, \end{aligned}$$

where the inequalities

$$\liminf_{t\to\infty} \quad \mathbb{E}_{\mathbb{P}^\star}\left\{\boldsymbol{x}_t^\top Q\boldsymbol{x}_t + \boldsymbol{u}_t^\top R\boldsymbol{u}_t\right\} \leq J_\infty(\pi) \leq \limsup_{t\to\infty} \quad \mathbb{E}_{\mathbb{P}^\star}\left\{\boldsymbol{x}_t^\top Q\boldsymbol{x}_t + \boldsymbol{u}_t^\top R\boldsymbol{u}_t\right\} \tag{22}$$

follow immediately from the definition of *limit inferior* and *limit superior*, respectively. Furthermore, the following can be asserted about the stationary behaviour of the stochastic process $\boldsymbol{x}_t$ when the controller $\pi \in \Pi_\infty$ stabilizes $\mathcal{S}$ asymptotically:

Theorem IV.1 (Steady state behaviour [21, Theorem 6.23]). *Let the discrete-time stochastic process $\boldsymbol{x}_t$ be the solution of the stochastic difference equation*

$$\boldsymbol{x}_{t+1} = \bar{A}\boldsymbol{x}_t + \bar{C}\boldsymbol{w}_t,$$

$$\boldsymbol{x}_{t_0} = x_0$$

where $\bar{A} \in \mathbb{R}^{n\times n}$, $\bar{C} \in \mathbb{R}^{n\times d}$ and $\boldsymbol{w}_t$ has zero mean and satisfies Assumption III.1. If $\bar{A}$ is strictly stable and $t_0 \to -\infty$, then the covariance

$$C_{xx}(s,t) := \mathbb{E}_{\mathbb{P}^\star}\left\{[\boldsymbol{x}_s - \mathbb{E}_{\mathbb{P}^\star}\{\boldsymbol{x}_s\}]\cdot[\boldsymbol{x}_t - \mathbb{E}_{\mathbb{P}^\star}\{\boldsymbol{x}_t\}]^\top\right\}$$

of the state process tends to an asymptotic value that depends only on the difference $s-t$. The asymptotic variance matrix $P_\infty := \lim_{t_0\to-\infty} C_{xx}(0,0)$ exists and is the unique solution of the discrete Lyapunov equation

$$P_\infty = R_{xx}(0) = \bar{A}P_\infty\bar{A}^\top + \bar{C}\bar{C}^\top.$$

In light of Theorem IV.1, we have that the inequalities in (22) are tight whenever Π_∞ is restricted to contain only exponentially stabilizing linear time-invariant control policies. Hence, we consider the following control problem

$$\begin{aligned}
\inf_{\pi\in\Pi_\infty} \quad & \lim_{t\to\infty} \quad \mathbb{E}_{\mathbb{P}^\star}\left\{\boldsymbol{x}_t^\top Q\boldsymbol{x}_t + \boldsymbol{u}_t^\top R\boldsymbol{u}_t\right\} \\
\text{s.t.} \quad & \boldsymbol{x}_{t+1} = A\boldsymbol{x}_t + B\boldsymbol{u}_t + C\boldsymbol{w}_t, \\
& \limsup_{t\to\infty} \ \sup_{\mathbb{P}\in\mathcal{P}_\infty} \mathbb{P}\text{-CVaR}_\epsilon\left(L_0\left(\boldsymbol{x}_t\right)\right) \leq 0.
\end{aligned} \tag{$\mathcal{R}_\infty$}$$

We assume throughout that the pairs $(Q^{1/2}, A)$ and (C, A) are observable and that the pair (A, B) is stabilisable, which is sufficient to guarantee the existence of linear time-invariant exponentially stabilising control policies. The assumption $\boldsymbol{x}_0 = 0$ is without loss of generality, i.e. the limit of the covariance matrix of the state is independent of $\boldsymbol{x}_0$ as indicated by Theorem IV.1. Under the observability and stabilisability assumptions made, the cost function is independent of the distribution of $\boldsymbol{x}_0$, as are the optimal control policies.

We restrict attention to linear control strategies for problem $\mathcal{R}_\infty$, for the same reasons mentioned in Section III. It is well known that such a restriction causes no loss of optimality when constraint (21) is disregarded. The equivalence (12) gives a probabilistic interpretation to what otherwise could be considered an *ad hoc* covariance constraint [33]. We next show that the equivalence (12) also implies that the optimal linear feedback law for problem $\mathcal{R}_\infty$ has an order which equals the number of states n of system $\mathcal{S}$, and is the combination of a Kalman filter and a static feedback gain.

Theorem IV.2 (Optimal linear feedback law)**.** *The optimal linear feedback law $\pi^\star$ of problem $\mathcal{R}_\infty$ consists of a linear estimator-controller pair (S, K) and hence is of the form*

$$\pi^\star : \begin{cases} \hat{\boldsymbol{x}}_{t+1} = (A + BK)\hat{\boldsymbol{x}}_t + S\left(\boldsymbol{y}_t - D\hat{\boldsymbol{x}}_t\right) \\ \quad \boldsymbol{u}_t = K\hat{\boldsymbol{x}}_t, \end{cases} \tag{23}$$

with $S := AYD^\top \left(DYD^\top + EE^\top\right)^{-1}$. The matrix Y is the unique positive definite solution of the discrete algebraic Riccati equation

$$Y = A\left(Y - YD^\top \left(DYD^\top + EE^\top\right)^{-1} DY\right) A^\top + CC^\top,$$

which can be solved efficiently [1]. The static feedback matrix is given by $K = Z^\star(P^\star)^{-1}$, where $P^\star \in \mathbb{S}^n_{++}$ and $Z^\star \in \mathbb{R}^{m\times n}$ can be found as the optimal solution of the SDP

$$\begin{aligned} \inf \quad & \operatorname{Tr} Q\,(Y + P) + \operatorname{Tr} RX \\ \text{s.t.} \quad & P \in \mathbb{S}^n_+,\ Z \in \mathbb{R}^{m\times n}, X \in \mathbb{S}^m_+ \\ & \begin{pmatrix} X & Z \\ Z^\top & P \end{pmatrix} \succeq 0,\ e^0 + \tfrac{1}{\epsilon}\operatorname{Tr}\left\{E_0\,(Y + P)\right\} \leq 0 \\ & \begin{pmatrix} P - APA^\top - BZA^\top - AZ^\top B^\top - \tilde{W} & BZ \\ Z^\top B^\top & P \end{pmatrix} \succeq 0 \end{aligned} \tag{24}$$

where $\tilde{W} := AYD^\top \left(DYD^\top + EE^\top\right)^{-1} DYA^\top$. Since (23) *can be decomposed into a Kalman estimator S and state feedback controller K, problem $\mathcal{R}_\infty$ satisfies a separation or certainty equivalence principle.*

Proof: See Appendix A. ■

The Kalman filter in Theorem IV.2 depends only on the process and measurement noise characteristics and is independent of the distributionally robust constraint (21) and cost function J_∞. Finding the optimal static feedback gain K requires only the solution of the tractable convex problem (24).

V. Wind turbine blade control design problem

To illustrate the method introduced in the preceding section, we consider a wind turbine control problem similar to the one introduced in [26]. As the size of wind turbines is increased for larger energy capture, they are subject to greater risks of fatigue failure and extreme loading events. Therefore, most large wind turbines today are equipped with pitch control for speed regulation, which can also be used for load alleviation.

However, these pitch actuators are slow and limited by the inertia of the blades. Hence, as in [26], we assume that the blades are equipped with an actively controlled flap. The control objective is to minimize actuation energy while keeping some measure of blade loading within specified bounds. The disturbance acting on the turbine blades is mostly due to atmospheric turbulence, for which little more than the frequency spectrum is known [10]. According to the standard design reference [24], atmospheric turbulence is typically treated as a Gaussian stochastic process defined by a standardized velocity spectrum. We follow the standard atmospheric turbulence model provided in [24], modulo the normality assumption which is not well supported in reality. Hence, this is a natural setting in which the ideas developed in this paper are of practical interest.

We start by briefly stating how a linear model of a wind turbine blade can be obtained. The modelling technique used here is by no means the only one possible, but results in a modest sized plant model of only ten states. An alternative technique using classical vortex-panel methods [19] to get higher fidelity, but still linear, models is presented in [26]. Since the disturbance behaviour is an important aspect of our model, we introduce it separately following the introduction of the physical model of the wind turbine blades. Some numerical results are provided in the last part of this section.

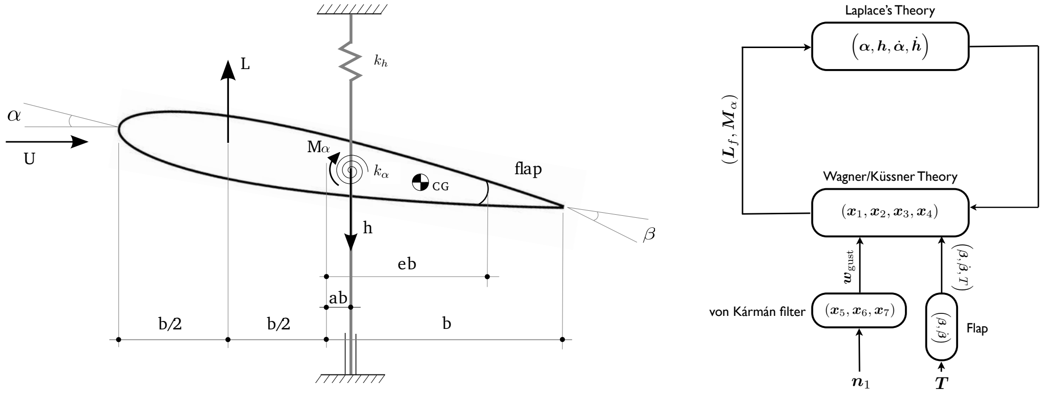

(a) (b)

Fig. 2. Figure 2(a) shows the geometry of the 2-DOF structural model. Figure 2(b) illustrates the connection of the different subcomponents in the overall wind turbine wing model. The overall model is linear continuous time invariant and has a modest size of 13 states, one endogenous and exogenous input $\boldsymbol{T}$ and $\boldsymbol{n}_1$, respectively.

A. Mathematical modelling

An aerofoil section with flap can be modelled using a simple two degree of freedom (2-DOF) plunge-pitch aerofoil, restrained by a pair of springs as shown in Figure 2(a). The two dimensional aerofoil represents a cross section of one of the flexible wind turbine blades. The spring constants are k_α which restrains pitch and k_h which restrains plunge motions. The free-stream velocity is denoted by U, pitch angle by α, plunge by h and flap deflection angle by β. The conventions taken for h is positive downwards, α is positive nose up about the elastic axis and β is positive flap down about the hinge. The lift L_f is positive upwards and the torque M_α about the elastic axis is positive when the aerofoil pitches nose up.

Structural modelling: The dynamic equations of motion in continuous time of the aerofoil can be obtained using Lagrange's equation [22] if structural damping is neglected. They can be written in state space form as

$$\begin{pmatrix} M & S_\alpha \\ S_\alpha & I_\alpha \end{pmatrix} \begin{pmatrix} \ddot{\boldsymbol{h}} \\ \ddot{\boldsymbol{\alpha}} \end{pmatrix} + \begin{pmatrix} k_h & 0 \\ 0 & k_\alpha \end{pmatrix} \begin{pmatrix} \boldsymbol{h} \\ \boldsymbol{\alpha} \end{pmatrix} = \begin{pmatrix} -\boldsymbol{L}_f \\ \boldsymbol{M}_\alpha \end{pmatrix}, \tag{25}$$

where M is the mass, S_α is the static moment of the aerofoil about the elastic axis and I_α is the aerofoil's moment of inertia about the elastic axis. The second derivatives of the pitch and plunge are denoted by $\ddot{\boldsymbol{\alpha}}$ and $\ddot{\boldsymbol{h}}$ respectively. The flaps are assumed to be rigid, i.e. they do not flex, and their mass is considered to be small, so any inertia effects from flap deflection can be ignored. As a result, flap deflection does not appear in equation (25) directly. The flap is actuated using torque control

$$m_f\ddot{\boldsymbol{\beta}} + r_f\dot{\boldsymbol{\beta}} + k_f\boldsymbol{\beta} = \boldsymbol{T},$$

where the friction factor $r_f \in \mathbb{R}_{++}$, flap inertia $m_f \ll M$, and flap spring k_f are given. The control input $\boldsymbol{T}$ is the torque forcing the wing flap. The power consumed by the flap is assumed to be proportional to the norm of $\dot{\boldsymbol{\beta}}$.

Linear unsteady aerodynamics: In order to derive a mathematical model for the aerodynamic disturbances, we must describe how the fluid flow around the aerofoil generates the lift $\boldsymbol{L}_f$ and torque $\boldsymbol{M}_\alpha$. An analytical method for describing 2D unsteady aerodynamics in the presence of transversal gust flow is provided by Wagner [30] and Küssner [20]. The closed form expression for the lift $\boldsymbol{L}_f$ and torque $\boldsymbol{M}_\alpha$ acting on the aerofoil was obtained by considering a 2D flat plate[5] with the additional flap undergoing harmonic motion in an inviscid and incompressible flow with flat wake at fixed free-stream velocity U and transversal gust flow at speed $\boldsymbol{w}_{\text{gust}}$.

The total lift $\boldsymbol{L}_f$ and total torque $\boldsymbol{M}_\alpha$ in this framework are given as the superposition of the lift caused by the longitudinal flow at speed U and transversal flow at speed $\boldsymbol{w}_{\text{gust}}$. The lift and torque caused by the longitudinal flow are determined using the Wagner filter [30]. Superimposing the lift and torque generated by the transversal gust flow $\boldsymbol{w}_{\text{gust}}$ determined using the Küssner filter [20], gives us an expression for the total lift $\boldsymbol{L}_f$ and total torque $\boldsymbol{M}_\alpha$. Both the Wagner and Küssner filter have a second order state space representation, resulting in four additional states $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_4)$, as shown in Figure 2(b).

Disturbance model: The majority of the disturbance acting upon the wind turbine blades is a direct result of atmospheric turbulence. Most commonly, atmospheric turbulence is represented as the convolution of (Gaussian) white noise through a linear time-invariant (LTI) shaping filter, usually referred to as a *von Kármán* filter, see [10], [24]. Hence

$$\boldsymbol{w}_{\text{gust}} := \mathcal{H} \cdot \boldsymbol{n}_1,$$

[5] The flat plate model is parametrised by a half chord length of b, a distance from mid-chord to elastic axis denoted by ab and a distance from mid-chord to hinge denoted by eb, see Figure 2(a).

where $\boldsymbol{n}_1$ is Gaussian white noise and $\mathcal{H}$ the *von Kármán* filter, which we choose to be a proper stable rational filter as in [10] with state space representation

$$\left[\begin{array}{ccc|c} -7.701 & -7.008 & -1.404 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \hline 1.447 & 7.022 & 1.533 & 0 \end{array}\right]. \tag{$\mathcal{H}$}$$

It is clear that the Gaussian assumption made on $\boldsymbol{w}_{\text{gust}}$ is unlikely to be fulfilled in practice, hence we assume only that $\boldsymbol{n}_1$ is a scalar white w.s.s. noise process, i.e. $\mathbb{E}\left\{\boldsymbol{n}_1^2(t)\right\} = 1$ and thus not necessarily Gaussian. Hence, in practice we need only estimate the power spectrum of the atmospheric turbulence $\boldsymbol{w}_{\text{gust}}$, e.g. from historical data.

The overall system of the wind turbine blade model with additional flap is a linear continuous time invariant system with 13 states, 6 states for the structural model, 4 states for the flow model and 3 states for the disturbance model. The overall model has one endogenous input $\boldsymbol{T}$ and one exogenous input $\boldsymbol{n}_1$ as shown in Figure 2(b). We assume that the states $\boldsymbol{\alpha}$ and $\boldsymbol{h}$ are measured with negligible measurement noise, i.e.

$$\boldsymbol{y} = \begin{pmatrix} \boldsymbol{\alpha} \\ \boldsymbol{h} \end{pmatrix} + \delta \boldsymbol{n}_2,$$

where $\boldsymbol{n}_2$ is a zero mean white noise signal with unit covariance matrix, uncorrelated with $\boldsymbol{n}_1$. To fit in the framework provided in the paper, we discretize the continuous time model using the zero order hold method at sampling frequency $f_s = 100$ Hz which captures most of the salient plant dynamics for the model parameters we have selected.

B. Numerical results

A natural control design criterion in this setting is to ensure that the vector $(\dot{\boldsymbol{\alpha}}, \dot{\boldsymbol{h}})$ is kept small in order to bound the fatigue stress, usually caused by high variance dynamic loads. In addition we would like extreme static loading events to be rare, corresponding to the requirement that the deformation vector $(\boldsymbol{\alpha}, \boldsymbol{h})$ remains close to zero. We express these two design criteria

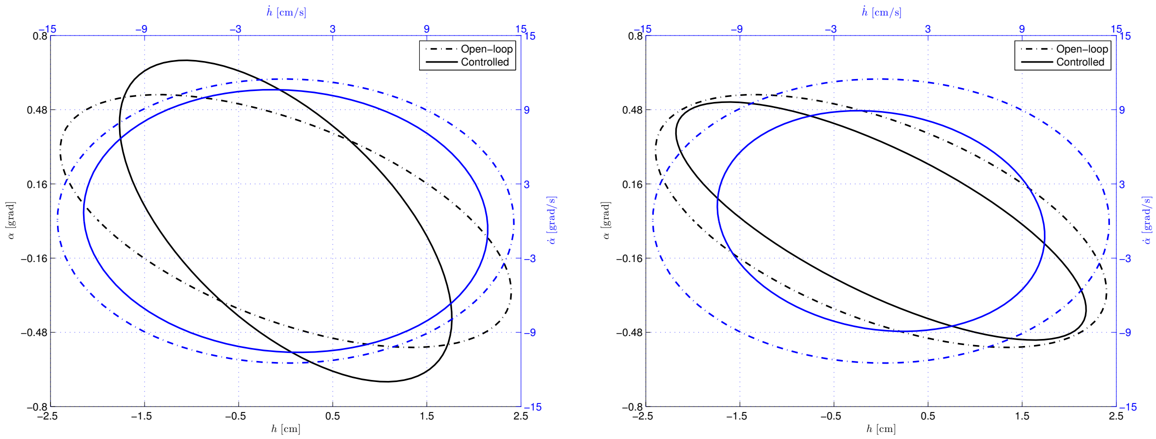

(a) (b)

Fig. 3. Figure 3(a) shows the variance of the vectors $(\boldsymbol{\alpha}, \boldsymbol{h})$ and $(\dot{\boldsymbol{\alpha}}, \dot{\boldsymbol{h}})$ when uncontrolled and with the optimal controller according to Section IV, as the sets $\left\{x \in \mathbb{R}^2 \;\middle|\; x^\top \Sigma^{-1} x - 1 \leq 0\right\}$ with Σ the respective covariance matrix. Similarly, Figure 3(b) shows the variance of the vectors $(\boldsymbol{\alpha}, \boldsymbol{h})$ and $(\dot{\boldsymbol{\alpha}}, \dot{\boldsymbol{h}})$ when uncontrolled, and with the standard LQR controller $K_{\mathrm{LQR}}(0.1)$.

respectively as

$$\limsup_{t \to \infty} \inf_{\mathbb{P} \in \mathcal{P}_\infty} \mathbb{P}\left\{(\dot{\boldsymbol{\alpha}}(t), \dot{\boldsymbol{h}}(t)) \in \mathcal{B}_2[55]\right\} \geq 1 - \epsilon, \tag{26}$$

$$\limsup_{t \to \infty} \inf_{\mathbb{P} \in \mathcal{P}_\infty} \mathbb{P}\left\{(\boldsymbol{\alpha}(t), \boldsymbol{h}(t)) \in \mathcal{B}_2[6]\right\} \geq 1 - \epsilon, \tag{27}$$

where $\epsilon = 0.1$, and $\mathcal{B}_n[r]$ denotes a closed ball in $\mathbb{R}^n$ of radius r around the origin. The natural control objective in this setting is to minimize the expected actuation power usage. We express this by taking as a cost function:

$$J(\pi) = \limsup_{t \to \infty} \sup_{\mathbb{P} \in \mathcal{P}_\infty} \mathbb{E}_{\mathbb{P}}\left\{\dot{\boldsymbol{\beta}}^2(t)\right\},$$

which must be minimized subject to the fatigue and loading constraints (26) and (27) respectively. Using the method described in Section IV, the optimal linear time invariant controller can be computed efficiently. Although it should be noted that in Theorem IV.2 only one probability constraint is considered, the generalisation to the case of finitely many constraints of type (21) is straightforward and omitted here. The difference between the variance of the vectors $(\boldsymbol{\alpha}, \boldsymbol{h})$ and $(\dot{\boldsymbol{\alpha}}, \dot{\boldsymbol{h}})$, when uncontrolled or controlled with the synthesized controller $K^\star$, is visualized in Figure 3(a).

Control	J	$(\boldsymbol{\alpha}, \boldsymbol{h}) \notin \mathcal{B}_2[6]$	$(\dot{\boldsymbol{\alpha}}, \dot{\boldsymbol{h}}) \notin \mathcal{B}_2[55]$
Uncontrolled	0	0.16	0.12
$K^\star$	82	0.10	0.10
$K_{\text{LQR}}(0.43)$	82	0.16	0.09
$K_{\text{LQR}}(0.1)$	425	0.15	0.07
$K_{\text{LQR}}(3.2 \times 10^{-3})$	3730	0.10	0.05

TABLE I

NUMERICAL RESULTS FOR THE WIND TURBINE BLADE CONTROL PROBLEM. THE THIRD AND FOURTH COLUMN SHOW THE WORST-CASE PROBABILITY THAT $(\boldsymbol{\alpha}, \boldsymbol{h}) \notin \mathcal{B}_2[6]$ AND $(\dot{\boldsymbol{\alpha}}, \dot{\boldsymbol{h}}) \notin \mathcal{B}_2[55]$, RESPECTIVELY.

We compare this controller to the standard $\mathcal{H}_2$-optimal controller found by tuning the cost function

$$\limsup_{t\to\infty} \quad \sup_{\mathbb{P}\in\mathcal{P}_\infty} \mathbb{E}_{\mathbb{P}} \left\{ \gamma \dot{\boldsymbol{\beta}}^2(t) + \boldsymbol{\alpha}^2(t) + \dot{\boldsymbol{\alpha}}^2(t) + \boldsymbol{h}^2(t) + \dot{\boldsymbol{h}}^2(t) \right\},$$

which weighs the actuation energy versus the size of the states $(\boldsymbol{\alpha}, \dot{\boldsymbol{\alpha}}, \boldsymbol{h}, \dot{\boldsymbol{h}})$, according to the tuning factor γ. A naïve method of designing a controller is to tune γ such that the closed loop system satisfies the fatigue (26) and loading (27) constraints.

We compare in Table V-B the cost of the optimal controller $K^\star$ and three naïvely tuned controllers $K_{\text{LQR}}(\gamma_i)$. First it is noted that when uncontrolled, the control cost is zero. However, since $\epsilon = 0.1$ both design specifications (26) and (27) are violated. The optimal controller $K^\star$ has satisfied (26) and (27) exactly with no conservatism and relatively low cost. The LQR controller $K_{\text{LQR}}(0.43)$ has the same cost as $K^\star$ but does not satisfy the constraints. The other LQR controllers either violate one of the constraints or have a massive cost compared to $K^\star$. The difference between the variance of the vectors $(\boldsymbol{\alpha}, \boldsymbol{h})$ and $(\dot{\boldsymbol{\alpha}}, \dot{\boldsymbol{h}})$, when uncontrolled or controlled with the controller $K_{\text{LQR}}(0.1)$, is visualized in Figure 3(b).

It can be seen from this example that the methodology of Section IV provides an easy procedure to design controllers that handle constraints of the type (26) and (27). Again we point out that, by dropping the Gaussian assumption on the stochastic process $(\boldsymbol{n}_1, \boldsymbol{n}_2)$, an assumption which in reality can not be justified anyway, the distributionally robust constraint formulation both makes practical sense and leads to a computationally tractable formulation.

VI. Conclusion

We investigate constrained control problems for stochastic linear systems when faced with the problem of only having limited information regarding the disturbance process, i.e. knowing only the first two moments of the disturbance distribution. We propose the use of distributionally robust chance and CVaR constraints to express constraint specifications when faced with distributional ambiguity.

These distributionally robust constrained formulations are subsequently used as control design specifications in both a finite horizon optimal control problem, and in an average cost optimal infinite horizon control problem.

We argue that these types of constraint formulations are practically meaningful and computationally tractable in the proposed finite and infinite horizon control design problems. The efficacy of the proposed formulation is illustrated for a wind turbine blade control design case study where flexibility issues play an important role and in which the distributionally robust framework makes practical sense.

VII. Acknowledgements

We would like to thank Bing Feng Ng and Rafael Palacios for introducing the gust load alleviation control problem and providing us with an appropriate model. The first author gratefully acknowledges the financial support from the Marie Curie FP7-ITN project “Energy savings from smart operation of electrical, process and mechanical equipment – ENERGY-SMARTOPS”, Contract No: PITN-GA-2010-264940. The second author’s work was supported by the Engineering and Physical Sciences Research Council (EP/I014640/1). The third author’s work was supported by the EU FP7 project EMBOCON (ICT-2009-4-248940).

APPENDIX

Define the matrices $\mathcal{B} \in \mathbb{R}^{N_x \times N_u}$, $\mathcal{C} \in \mathbb{R}^{N_x \times N_w}$, $\mathcal{D} \in \mathbb{R}^{N_y \times N_u}$ and $\mathcal{E} \in \mathbb{R}^{N_y \times N_w}$ as follows

$$\mathcal{B} := \begin{pmatrix} 0 & & & & \\ B & 0 & & & \\ AB & B & 0 & & \\ \vdots & & \ddots & \ddots & \\ \vdots & & & B & 0 \\ A^{N-1}B & A^{N-2}B & \dots & AB & B \end{pmatrix} =: \begin{pmatrix} \mathcal{B}_0 \\ \mathcal{B}_1 \\ \mathcal{B}_2 \\ \vdots \\ \vdots \\ \mathcal{B}_N \end{pmatrix}, \quad \mathcal{D} := \begin{pmatrix} 0 & & & \\ D & 0 & & \\ & D & 0 & \\ & & \ddots & \ddots \\ & & & D \quad 0 \end{pmatrix} =: \begin{pmatrix} \mathcal{D}_0 \\ \mathcal{D}_1 \\ \mathcal{D}_2 \\ \vdots \\ \mathcal{D}_{N-1} \end{pmatrix},$$

$$\mathcal{C} := \begin{pmatrix} x_0 & & & \\ Ax_0 & C & & \\ A^2x_0 & AC & C & \\ \vdots & & \ddots & \ddots \\ A^N x_0 & A^{N-1}C & \dots & AC \quad C \end{pmatrix} =: \begin{pmatrix} \mathcal{C}_0 \\ \mathcal{C}_1 \\ \mathcal{C}_2 \\ \vdots \\ \mathcal{C}_N \end{pmatrix}, \qquad \mathcal{E} := \begin{pmatrix} 1 & & & \\ & E & & \\ & & E & \\ & & & \ddots \\ & & & & E \end{pmatrix} =: \begin{pmatrix} \mathcal{E}_0 \\ \mathcal{E}_1 \\ \mathcal{E}_2 \\ \vdots \\ \mathcal{E}_{N-1} \end{pmatrix},$$

where x_0 is the initial state of system $\mathcal{S}$, and $N_x := (N+1)n$, $N_u := Nm$, $N_w := Nd + 1$ and $N_y = rN$.

The following general theorem establishes that distributionally robust constraints are invariant under linear projections. It plays an important role in the proof of Theorem IV.2.

Theorem A.1 (General projection property)**.** *Let $\boldsymbol{x}$ and $\boldsymbol{w}$ be random vectors valued in $\mathbb{R}^n$ and $\mathbb{R}^d$, respectively, and define the ambiguity sets*

$$\mathcal{P}_x := \left\{ \mathbb{P} \in \mathcal{P}_0 \;\middle|\; \mathbb{E}_{\mathbb{P}}\left\{(\boldsymbol{x}^\top, 1)^\top(\boldsymbol{x}^\top, 1)\right\} = M_x \right\}$$

and

$$\mathcal{P}_w := \left\{ \mathbb{P} \in \mathcal{P}_0 \;\middle|\; \mathbb{E}_{\mathbb{P}}\left\{(\boldsymbol{w}^\top, 1)^\top(\boldsymbol{w}^\top, 1)\right\} = M_w \right\},$$

where $M_x \in \mathbb{S}_+^{n+1}$ and $M_w \in \mathbb{S}_+^{d+1}$ are related through $M_x = \left(\begin{smallmatrix} X & 0 \\ 0^\top & 1 \end{smallmatrix}\right) M_w \left(\begin{smallmatrix} X & 0 \\ 0^\top & 1 \end{smallmatrix}\right)^\top$ for some $X \in \mathbb{R}^{n \times d}$. Then, for any Borel measurable function $f : \mathbb{R}^n \to \mathbb{R}$ that admits a quadratic minorant we have

$$\inf_{\mathbb{P} \in \mathcal{P}_x} \mathbb{E}_{\mathbb{P}}(f(\boldsymbol{x})) = \inf_{\mathbb{P} \in \mathcal{P}_w} \mathbb{E}_{\mathbb{P}}(f(X\boldsymbol{w})).$$

Proof: This is an immediate consequence of [32, Theorem 1]. ■

Proof of Theorem III.1

The proof follows by applying the tractability result in Theorem II.1 to the constraints

$$\sup_{\mathbb{P} \in \mathcal{P}_\infty} \mathbb{P}\text{-}\mathbf{CVaR}_\epsilon(L(\overbrace{(\mathcal{B}_t U + \mathcal{C}_t)\,\boldsymbol{w}}^{\boldsymbol{x}_t}; \alpha)) \leq 0.$$

Explicitly writing out the quadratic form in the preceding inequality as

$$\sup_{\mathbb{P}\in\mathcal{P}_\infty} \mathbb{P}\text{-CVaR}_\epsilon \left(\max_i \alpha_i \boldsymbol{w}^\top (\mathcal{B}_t U + \mathcal{C}_t)^\top E_i (\mathcal{B}_t U + \mathcal{C}_t) \boldsymbol{w} + 2\alpha_i e_i^\top (\mathcal{B}_t U + \mathcal{C}_t) \boldsymbol{w} + \alpha_i e_i^0 \right) \leq 0$$

yields a matrix inequality with quadratic terms in the variable U:

$$\exists \beta_t \in \mathbb{R},\ X_t \in \mathbb{S}_+^{Nd+2} : \begin{cases} \beta_t + \frac{1}{\epsilon}\operatorname{Tr}\{M_w X_t\} \leq 0 \\ X_t \succeq \begin{pmatrix} \alpha_i (\mathcal{B}_t U + \mathcal{C}_t)^\top E_i (\mathcal{B}_t U + \mathcal{C}_t) & (\mathcal{B}_t U + \mathcal{C}_t)^\top e_i \alpha_i \\ \alpha_i e_i^\top (\mathcal{B}_t U + \mathcal{C}_t) & e_i^0 \alpha_i - \beta_t \end{pmatrix}, \forall i \in \{1, \dots, I\} \end{cases}$$

The final result claimed in the theorem is than found by applying Schur complements, and rewriting the quadratic matrix inequality as two LMIs using the additional variables $P_t^i \in \mathbb{S}_+^{Nd+1}$.

Proof of Theorem IV.2

The feedback policies π in problem $\mathcal{R}_\infty$ are restricted to be linear and causal, i.e. $\exists X_t \in \mathbb{R}^{n\times td} : \boldsymbol{x}_t = X_t T_t \boldsymbol{w}$, where $T_t : w \to (w_0^\top, \dots, w_{t-1}^\top)^\top$ is a truncation. We have then according to Theorem A.1,

$$\sup_{\mathbb{P}\in\mathcal{P}_w} \mathbb{P}\text{-CVaR}_\epsilon \left(L_0 \left(X_t T_t \boldsymbol{w} \right) \right) \leq 0 \iff \sup_{\mathbb{P}\in\mathcal{P}_{x_t}} \mathbb{P}\text{-CVaR}_\epsilon \left(L_0(\boldsymbol{x}_t) \right) \leq 0,$$

where $M_{x_t} = \operatorname{diag}(X_t T_t, 1) \cdot M_w \cdot \operatorname{diag}(X_t T_t, 1)^\top$. Now according to Corollary II.1, we have that the limit as $t \to \infty$ of the preceding inequality is equivalent to

$$\lim_{t\to\infty} e_0 + \frac{1}{\epsilon}\operatorname{Tr}\left\{ E_0^{1/2} \mathbb{E}_{\mathbb{P}^\star}\left\{ \boldsymbol{x}_t \boldsymbol{x}_t^\top \right\} E_0^{1/2} \right\} \leq 0,$$

where Theorem IV.1 guarantees that the limit exists. The objective function can be written in the form of a standard $\mathcal{H}_2$-problem,

$$J_{\text{lqr}} = \lim_{t\to\infty} \operatorname{Tr}\left\{ Q^{1/2} \mathbb{E}_{\mathbb{P}^\star}\left\{ \boldsymbol{x}_t \boldsymbol{x}_t^\top \right\} Q^{1/2} + R^{1/2} \mathbb{E}_{\mathbb{P}^\star}\left\{ \boldsymbol{u}_t \boldsymbol{u}_t^\top \right\} R^{1/2} \right\},$$

using the fact that the expectation operator is linear and the trace identity $\operatorname{Tr}\{AB\} = \operatorname{Tr}\{BA\}$. Hence, when restricted to linear control strategies, problem $\mathcal{R}_\infty$ reduces to

$$\begin{aligned} \inf_\pi \quad & \lim_{t\to\infty} \quad \sup_{\mathbb{P}\in\mathcal{P}_\infty} \mathbb{E}_{\mathbb{P}}\left\{ \boldsymbol{x}_t^\top Q \boldsymbol{x}_t + \boldsymbol{u}_t^\top R \boldsymbol{u}_t \right\} \\ \text{s.t.} \quad & \boldsymbol{x}_{t+1} = A\boldsymbol{x}_t + B\boldsymbol{u}_t + C\boldsymbol{w}_t, \\ & \lim_{t\to\infty} \quad \sup_{\mathbb{P}\in\mathcal{P}_\infty} \operatorname{Tr}\left\{ E^{1/2} \mathbb{E}_{\mathbb{P}}\left\{ \boldsymbol{x}_t \boldsymbol{x}_t^\top \right\} E^{1/2} \right\} \leq -e_0 \epsilon. \end{aligned}$$

However, the last problem is an instance of a standard multi-criterion $\mathcal{H}_2$-problem, see [5, §12.2.1]. The fact that the optimal control law is of the form (23) is a result of the fact that it solves an $\mathcal{H}_2$-problem with a different cost measure, i.e. there exists an unconstrained $\mathcal{H}_2$-problem with state and input penalty matrices $\tilde{Q}$, $\tilde{R}$ for which the solution satisfies the omitted trace constraint [5, §6.5.1]. The fact that K can be found as the solution to an SDP can be found in [6], and essentially follows from standard LMI manipulations.

REFERENCES

[1] W.F. Arnold III and A.J. Laub. Generalized eigenproblem algorithms and software for algebraic Riccati equations. *Proceedings of the IEEE*, 72(12):1746–1754, 1984.

[2] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, 2009.

[3] D.P. Bertsekas and I.B. Rhodes. On the minimax reachability of target sets and target tubes. *Automatica*, 7(2):233–247, 1971.

[4] F. Blanchini. Set invariance in control. *Automatica*, 35(11):1747–1767, 1999.

[5] S.P. Boyd and C.H. Barratt. *Linear Controller Design: Limits of Performance*. Prentice Hall, 1991.

[6] S.P. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in Sstem and Control Theory*. Society for Industrial and Applied Mathematics, 1987.

[7] S.P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[8] G.C. Calafiore. Random convex programs. *SIAM Journal on Optimization*, 20(6):3427–3464, 2010.

[9] G.C. Calafiore and M.C. Campi. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006.

[10] C.W. Campbell. Monte Carlo turbulence simulation using rational approximations to von Kármán spectra. *AIAA Journal*, 24(1):62–66, 1986.

[11] M. Cannon, B. Kouvaritakis, and X. Wu. Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty. *IEEE Transactions on Automatic Control*, 54(7):1626–1632, 2009.

[12] A. Charnes and W.W. Cooper. Chance-constrained programming. *Management Science*, 6(1):73–79, 1959.

[13] W. Chen, M. Sim, J. Sun, and C.-P. Teo. From CVaR to uncertainty set: Implications in joint chance-constrained optimization. *Operations Research*, 58(2):470–485, 2010.

[14] H. Föllmer and A. Schied. *Stochastic Finance. An Introduction in Discrete Time*. Number 27 in de Gruyter Studies in Mathematics. de Gruyter, 2004.

[15] P.J. Goulart and E.C. Kerrigan. Output feedback receding horizon control of constrained systems. *International Journal of Control*, 80(1):8–20, 2007.

[16] P.J. Goulart, E.C. Kerrigan, and J.M. Maciejowski. Optimization over state feedback policies for robust control with constraints. *Automatica*, 42(4):523–533, 2006.

[17] M.J. Hadjiyiannis, P.J. Goulart, and D. Kuhn. An efficient method to estimate the suboptimality of affine controllers. *IEEE Transactions on Automatic Control*, 56(12):2841–2853, 2011.

[18] D.A. Iancu, M. Sharma, and M. Sviridenko. Supermodularity and affine policies in dynamic robust optimization. 2013. To appear in *Operations Research*.

[19] J. Katz and A. Plotkin. *Low-Speed Aerodynamics*. Number 13 in Cambridge aerospace series. Cambridge University Press, 2001.

[20] H.G. Küssner. Zusammenfassender Bericht über den instationären Auftrieb von Flügeln. *Luftfahrtforschung*, 13(12):410–424, 1936.

[21] H. Kwakernaak and R. Sivan. *Linear Optimal Control Systems*, volume 172. Wiley-Interscience, 1972.

[22] J.G. Leishman. *Principles of Helicopter Aerodynamics*. Cambridge University Press, 2006.

[23] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[24] U.S. Military. U.S. Military Handbook MIL-HDBK-1797. Technical report, 1997.

[25] A. Nemirovski and A. Shapiro. Convex approximations of chance constrained programs. *SIAM Journal on Optimization*, 17(4):969–996, 2006.

[26] B.F. Ng, R. Palacios, J.M.R. Graham, and E.C. Kerrigan. Robust control synthesis for gust load alleviation from large aeroelastic models with relaxation of spatial discretisation. *Proceedings of the EWEA*, 2012.

[27] B.P. Van Parys, P.J. Goulart, and M. Morari. Infinite horizon performance bounds for uncertain constrained systems. *IEEE Transactions on Automatic Control*, submitted.

[28] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, 2002.

[29] R.T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42, 2000.

[30] H. Wagner. Über die Entstehung des dynamischen Auftriebes von Tragflügeln. *Zeitschrift für Angewandte Mathematik und Mechanik*, 5(1):17–35, 1925.

[31] H.S. Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control*, 6(1):131–147, 1968.

[32] Y.L. Yu, Y. Li, D. Schuurmans, and C. Szepesvári. A general projection property for distribution families. *Proceedings of the 23rd Annual Conference on Advances in Neural Information Processing Systems*, pages 2232–2240, 2009.

[33] G. Zhu, M.A. Rotea, and R. Skelton. A convergent algorithm for the output covariance constraint control problem. *SIAM Journal on Control and Optimization*, 35(1):341–361, 1997.

[34] S. Zhu and M. Fukushima. Worst-case conditional value-at-risk with application to robust portfolio management. *Operations Research*, 57(5):1155–1168, 2009.

[35] S. Zymler, D. Kuhn, and B. Rustem. Distributionally robust joint chance constraints with second-order moment information. *Mathematical Programming*, 137(1-2):167–198, 2013.