Optimality of Affine Policies in Multi-stage Robust Optimization

Dimitris Bertsimas, Dan A. Iancu and Pablo A. Parrilo

***Abstract*— In this paper, we prove the optimality of disturbance-affine control policies in the context of one-dimensional, box-constrained, multi-stage robust optimization. Our results cover the finite horizon case, with minimax (worst-case) objective, and convex state costs plus linear control costs. Our proof methodology, based on techniques from polyhedral geometry, is elegant and conceptually simple, and entails efficient algorithms for the case of piecewise affine state costs, when computing the optimal affine policies can be done by solving a single linear program.**

I. INTRODUCTION

Multi-stage optimization problems under uncertainty have been prevalent in numerous fields of science and engineering, and have elicited interest from diverse research communities, on both a theoretical and a practical level. Several solution approaches have been proposed, with various degrees of generality, tractability, and performance guarantees. Some of the most successful ones include exact and approximate dynamic programming, stochastic programming, sampling-based methods, and, more recently, robust and adaptive optimization, which is the focus of the present paper.

The topics of robust optimization and robust control have been studied, under different names, by a variety of academic groups, mostly in control theory (see [1], [2], and references therein) and operations research ([3], [4], [5]), with considerable effort put into justifying the assumptions and general modeling philosophy. As such, the goal of the current paper is not to *motivate* the use of robust (and, more generally, distribution-free) techniques. Rather, we take the modeling approach as a given, and investigate tractability and performance issues in the context of a certain class of optimization problems. More precisely, we are concerned with the following decision problem.

Problem 1.1: Consider a one-dimensional, discrete, linear, time-varying dynamical system,

$$x_{k+1} = \alpha_k \cdot x_k + \beta_k \cdot u_k + \gamma_k \cdot w_k,$$

where $\alpha_k, \beta_k, \gamma_k \neq 0$ *are known scalars, and the initial state* $x_1 \in \mathbb{R}$ *is specified. The random disturbances* w_k *are unknown, but bounded,*

$$w_k \in \mathscr{W}_k \stackrel{def}{=} [\underline{w}_k, \overline{w}_k].$$

D. Bertsimas is with Faculty of the Sloan School of Management and the Operations Research Center, Massachusetts Institute of Technology, 77 Massachusetts Avenue, E40-147 dbertsim@mit.edu

D. A. Iancu is with the Operations Research Center, Massachusetts Institute of Technology, 77 Massachusetts Avenue, E40-130 daniancu@mit.edu

P. A. Parrilo is with the Laboratory for Information and Decision Systems and the Faculty of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 77 Massachusetts Avenue, 32D-726 parrilo@mit.edu

We would like to find a sequence of bounded controllers,

$$u_k \in [L_k, U_k], \tag{1}$$

minimizing the following cost over a finite horizon $1, \ldots, T$*,*

$$J_1(x_1) \stackrel{def}{=} c_1 u_1 + \max_{w_1} \Big[h_1(x_2) + \ldots + c_k u_k + \max_{w_k} \Big[h_k(x_{k+1}) + \ldots + c_T u_T + \max_{w_T} \ h_T(x_{T+1}) \ \Big] \ldots \Big], \tag{2}$$

where the functions $h_k : \mathbb{R} \to \mathbb{R}$ *are convex and coercive, and* $c_k \geq 0$ *are fixed.*

The above problem can be seen as a game between the decision maker and nature, in which, at every stage k, the latter chooses a disturbance w_k maximizing the objective function, while the former chooses a constrained control action u_k minimizing the objective. Examples of such problems include the case of quadratic state costs, as well as norm-1 or norm-∞ costs, all of which have been studied extensively in the literature in the unconstrained case (see [1], [2]).

The solution to Problem 1.1 can be obtained using a "classical" Dynamic Programming formulation [6], in which the optimal policies $u_k^\star(x_k)$ and the optimal value functions $J_k^\star(x_k)$ are computed backwards in time, starting at the end of the planning horizon, $k = T$. The resulting policies are piecewise affine in the state x_k, and have properties that are well known and documented in the literature [7].

In the current paper, we study the performance of a new class of policies, where instead of regarding the controllers u_k as functions of the state x_k, one seeks direct parameterizations in the observed disturbances,

$$u_k : \mathscr{W}_1 \times \mathscr{W}_2 \times \cdots \times \mathscr{W}_{k-1} \to \mathbb{R}, \tag{3}$$

which are robustly feasible for constraint (1), i.e., $u_k(\boldsymbol{\xi}) \in [L_k, U_k]$, $\forall \boldsymbol{\xi} \in \mathscr{W}_1 \times \cdots \times \mathscr{W}_{k-1}$.

While such parameterizations require a state-space that is increasing in time, thereby potentially leading to more complicated optimization problems, the hope is that simpler functional forms, e.g. affine, might be sufficient for optimality. Such affine policies have a very compact representation (only the coefficients are needed), and, for certain classes of convex costs h_k, there may be efficient procedures available for computing them.

This approach is not new in the literature. It has been originally advocated in the context of stochastic programming [8], and then rediscovered in linear systems theory [9] and robust optimization ([10], [11]), with notable contributions from researchers in robust model predictive control and receding horizon control (see [7], [12], [13], and references therein). In all the papers, which usually deal with the more general case of multi-dimensional linear systems, the authors show how the reformulation can be done, and how the corresponding affine policies can be found by solving specific types of optimization problems, which vary from linear and quadratic ([11], [12]) to conic and semi-definite programs ([11], [9], [14], [15]). The first steps towards analyzing the properties of such parameterizations were made in [12], where the authors show that, under suitable conditions, the resulting affine parameterization has certain desirable system theoretic properties (stability and robust invariance). Another notable contribution was [16], where the authors prove that the class of affine disturbance feedback policies is equivalent to the class of affine state feedback policies with memory of prior states, thus subsuming the classes of open-loop and pre-stabilizing control policies. However, to the best of our knowledge, apart from these theoretical advances, there has been very little progress in proving results about the quality of the objective function values resulting from the use of such parameterizations.

Our main result, summarized in Theorem 3.1 of Section III, is that disturbance-affine policies *are optimal* for Problem 1.1. Furthermore, we prove that a certain affine relaxation of the state costs is possible, without any loss of optimality, which gives rise to very efficient algorithms for computing the optimal affine policies when the state costs h_k are piece-wise affine. To the best of our knowledge, this is the first result of its kind, and it provides intuition and motivation for the widespread advocation of such policies in both theory and applications. Our theoretical results are tight (if the conditions in the problem are slightly perturbed, then simple counterexamples for Theorem 3.1 can be found), and the proof of the theorem itself is atypical, consisting of a forward induction and making use of polyhedral geometry to construct the optimal affine policies.

The paper is organized as follows. Section II presents an overview of the Dynamic Programming formulation in state variable x_k, extracting the optimal policies $u_k^\star(x_k)$ and optimal value functions $J_k^\star(x_k)$, as well as some of their properties. Section III contains our main result, and briefly discusses some immediate extensions and computational implications. In Section IV, we introduce the constructive proof for building the affine control policies and the affine cost relaxations, and briefly discuss counterexamples that prevent a generalization of the results. Section V presents our conclusions and directions for future research.

A. NOTATION

Throughout the rest of the paper, the subscript k is used to denote time-dependency, and vector quantities are distinguished by bold-faced symbols. Since we seek policies parametrized directly in disturbances, we introduce $\boldsymbol{w}_k \stackrel{\text{def}}{=} (w_1,\ldots,w_{k-1})$ to denote the history of known disturbances in period k, and $\mathscr{H}_k \stackrel{\text{def}}{=} \mathscr{W}_1 \times \cdots \times \mathscr{W}_{k-1}$ to denote the corresponding uncertainty set.

A function q_k that depends affinely on variables $w_1,\ldots,w_{k-1}$ is denoted by $q_k(\boldsymbol{w}_k) \stackrel{\text{def}}{=} q_{k,0} + \boldsymbol{q}_k' \boldsymbol{w}_k$, where $\boldsymbol{q}_k$ is the vector of coefficients, and $'$ denotes the usual transpose operation. Optimal quantities have a $\star$ superscript, e.g., $J_k^\star$.

Since our exposition relies heavily on sets given by maps $\theta : \mathbb{R}^k \mapsto \mathbb{R}^2 (k \geq 2)$, in order to reduce the number of symbols, we denote the resulting coordinates in $\mathbb{R}^2$ by θ_1, θ_2, and use the following overloaded notation:

- $\theta_i(\boldsymbol{w})$ designates the value assigned by map θ to $\boldsymbol{w} \in \mathbb{R}^k$
- $\theta_i[\boldsymbol{v}]$ denotes the θ_i-coordinate of the point $\boldsymbol{v} \in \mathbb{R}^2$.

The different use of parentheses should remove any ambiguity from the notation (particularly in the case $k = 2$).

We use $\text{conv}(\cdot)$ and $\text{ext}(\cdot)$ to denote the convex hull and the set of extreme points, respectively.

II. DYNAMIC PROGRAMMING SOLUTION

As mentioned in the introduction, the solution to Problem 1.1 can be obtained using a "classical" Dynamic Programming (DP) formulation [6], in which the state is taken to be x_k, and the optimal policies $u_k^\star(x_k)$ and optimal value functions $J_k^\star(x_k)$ are computed starting at the end of the planning horizon, $k = T$, and moving backwards in time. In this section, we outline the DP solution for our problem, and state some of the key properties that are used throughout the rest of the paper. For completeness, a proof of the results is included in the journal version of our paper [17].

Since the constraints on the controls u_k and the bounds on the disturbances w_k are independent across time, we can restrict attention, without loss of generality[1], to a system with $\alpha_k = \beta_k = \gamma_k = 1$. With this simplification, the Bellman recursion for Problem 1.1 can be written as

$$J_k^\star(x_k) \stackrel{\text{def}}{=} \min_{L_k \leq u_k \leq U_k} \Big[c_k\, u_k + \max_{w_k \in \mathscr{W}_k} \big[h_k(x_k+u_k+w_k) + J_{k+1}^\star(x_k+u_k+w_k) \big] \Big],$$

where $J_{T+1}^\star(x) \equiv 0$. By defining:

$$y_k \stackrel{\text{def}}{=} x_k + u_k, \tag{4a}$$

$$g_k(y_k) \stackrel{\text{def}}{=} \max_{w_k \in \mathscr{W}_k} \big[\, h_k(y_k+w_k) + J_{k+1}^\star(y_k+w_k) \,\big], \tag{4b}$$

we obtain the following solution to the Bellman equation:

$$u_k^\star(x_k) = \begin{cases} U_k, & \text{if } x_k < \underline{y}_k - U_k \\ -x_k + y_k^\star, & \text{otherwise} \\ L_k, & \text{if } x_k > \overline{y}_k - L_k \end{cases} \tag{5a}$$

$$J_k^\star(x_k) = \begin{cases} c_k \cdot U_k + g_k(x_k+U_k), & \text{if } x_k < \underline{y}_k - U_k \\ c_k \cdot (y_k^\star - x_k) + g_k(y_k^\star), & \text{otherwise} \\ c_k \cdot L_k + g_k(x_k+L_k), & \text{if } x_k > \overline{y}_k - L_k\ , \end{cases} \tag{5b}$$

where $y_k^\star \in [\underline{y}_k, \overline{y}_k]$, and $[\underline{y}_k, \overline{y}_k]$ is the set of minimizers of the convex function $c_k \cdot y + g_k(y)$. In particular, we have that:

P1 The optimal control law $u_k^\star(x_k)$ is piecewise affine, with at most 3 pieces, continuous and non-increasing.

P2 The optimal value function, $J_k^\star(x_k)$, and the function $g_k(y_k)$ are convex.

[1] Such a system can always be obtained by the linear change of variables $\tilde{x}_k = \frac{x_k}{\prod_{i=1}^{k-1} \alpha_i}$, and by suitably scaling the bounds $L_k, U_k, \underline{w}_k, \overline{w}_k$.

III. OPTIMALITY OF AFFINE POLICIES IN DISTURBANCES

In this section, we introduce our main contribution, namely a proof that policies that are affine in the disturbances $\boldsymbol{w}_k$ are, in fact, optimal for Problem 1.1. Using the same notation as in Section II, and defining $J_{mM} = J_1^\star(x_1)$ as the overall optimal value (recall that x_1 is fixed), we summarize our main result in the following theorem:

Theorem 3.1: For every time step $k = 1, \ldots, T$, the following quantities exist:

- *an affine control policy, $q_k(\boldsymbol{w}_k)$*
- *an affine running cost, $z_k(\boldsymbol{w}_{k+1})$*

such that the following properties are obeyed:

$$L_k \le q_k(\boldsymbol{w}_k) \le U_k, \ \forall\, \boldsymbol{w}_k \in \mathscr{H}_k \tag{6a}$$

$$z_k(\boldsymbol{w}_{k+1}) \ge h_k\Big(x_1 + \sum_{t=1}^{k}\big(q_t(\boldsymbol{w}_t) + w_t\big)\Big), \ \forall\, \boldsymbol{w}_{k+1} \in \mathscr{H}_{k+1} \tag{6b}$$

$$J_{mM} = \max_{w_1,\ldots,w_k}\Bigg[\sum_{t=1}^{k}\big(c_t \cdot q_t(\boldsymbol{w}_t) + z_t(\boldsymbol{w}_{t+1})\big) + J_{k+1}^\star\Big(x_1 + \sum_{t=1}^{k}\big(q_t(\boldsymbol{w}_t)+w_t\big)\Big)\Bigg]. \tag{6c}$$

To understand the meaning of the claims, note that (6a) confirms that the affine policy $q_k(\boldsymbol{w}^k)$ is robustly feasible, i.e., it obeys the control constraints, for any realization of the disturbances. Equation (6b) states that the affine cost $z_k(\boldsymbol{w}^{k+1})$ is always larger than the convex state cost $h_k(x_{k+1})$, which would be incurred if the affine policies $q_t(\cdot), 1 \le t \le k$, were used. Equation (6c) guarantees that, despite using the (suboptimal) affine control law q_k, and incurring a (potentially larger) affine stage cost z_k, the overall objective function value J_{mM} is, in fact, not increased. This translates into the following two main results:

- *Existential result.* Affine policies $q_k(\boldsymbol{w}_k)$ are, in fact, optimal for Problem 1.1.
- *Computational result.* When the costs h_k are piecewise affine, the optimal affine policies can be computed by solving a single Linear Programming (LP) problem, with size polynomial in the problem input.

To see why the second result holds, note that if h_k is piecewise affine and convex, the original optimization problem can be written as a semi-infinite LP [18]. A typical constraint becomes bi-affine in the decision variables ($\boldsymbol{x}$) and uncertainties ($\boldsymbol{w}$), i.e., $\lambda_0(\boldsymbol{x}) + \sum_{t=1}^{T}\lambda_t(\boldsymbol{x}) \cdot w_t \le 0, \forall\, w_t \in \mathscr{W}_t, t = 1, \ldots, T$, where λ_i are affine in $\boldsymbol{x}$. It can be shown [10], [19] that such a condition is equivalent to:

$$\begin{cases} \lambda_0(\boldsymbol{x}) + \sum_{t=1}^{T}\Big(\lambda_t(\boldsymbol{x}) \cdot \frac{\underline{w}_t + \overline{w}_t}{2} + \frac{\overline{w}_t - \underline{w}_t}{2} \cdot \xi_t\Big) \le 0 \ , \\ -\xi_t \le \lambda_t(\boldsymbol{x}) \le \xi_t, \quad \forall t = 1, \ldots, T, \end{cases}$$

which are linear constraints in the decision variables $\boldsymbol{x}, \boldsymbol{\xi}$. Therefore, the problem of finding the optimal coefficients $\{q_{k,t}\}, \{z_{k,t}\}$ can be reformulated as an LP with $O\big(T^2 \cdot \max_k m_k\big)$ variables and $O\big(T^2 \cdot \max_k m_k\big)$ constraints, where m_k is the number of pieces in h_k.

We conclude our preliminary observations by noting that state constraints of the form $L_k^x \le x_k \le U_k^x$ could also be included in Problem 1.1. More precisely, if the mathematical problem including such constraints remains feasible[2], then disturbance-affine policies are still optimal. The reason is that a problem with convex stage costs h_k and state constraints $L_k^x \le x_k \le U_k^x$ is equivalent to a problem without any state constraints, but with modified, convex[3] stage costs $\tilde{h}_k \stackrel{\text{def}}{=} h_k + \mathbf{1}_{[L_{k+1}^x, U_{k+1}^x]}$ (where $\mathbf{1}_K$ is the indicator function of the set K), for which affine policies are optimal, by Theorem 3.1.

IV. PROOF OF MAIN THEOREM

In the current section, we present a sketch of the proof of Theorem 3.1. Our emphasis is on the intuition and key results that make the statements true. The interested reader is referred to [17] for full details.

Unlike most DP proofs, which utilize backward induction on the time-periods, we proceed with a *forward* induction. Section IV-A presents a test of the first step of the induction, and introduces a detailed analysis of the consequences of the induction hypothesis.

We then separate the completion of the inductive step into two parts. In the first part, discussed in Section IV-B, by exploiting the structure provided by the forward induction hypothesis, and making critical use of the properties of the optimal control law $u_k^\star$ and optimal value function $J_k^\star$, we introduce a candidate affine policy q_k, which can be proven to be robustly feasible, and preserving the min-max value of the overall problem, J_{mM}, when used in conjunction with the original, convex state costs, $h_k(x_{k+1})$.

Similarly, for the second part of the inductive step, treated in Section IV-C, by re-analyzing the feasible sets of the optimization problems resulting after the use of the affine policy q_k, we determine a candidate affine cost z_k, which can be shown to be always larger than the original convex state costs, h_k, but which, if incurred at time k, would leave the overall min-max value unchanged.

Section IV-D concludes our analysis by outlining several counterexamples that prevent an immediate extension of the result to more general cases.

A. Induction Hypothesis.

We first verify the induction at time $k = 1$. Since x_1 is fixed and $u_1^\star(\cdot)$ is always feasible, we take the affine control as the constant $q_1 \stackrel{\text{def}}{=} u_1^\star(x_1)$, so that (6a) is immediately obeyed. Furthermore, since $u_1^\star$ is optimal, the optimal objective value, $J_{mM} = J_1^\star(x_1)$, can be written as:

$$\begin{aligned} J_{mM} &= c_1 q_1 + g_1(x_1 + q_1) \stackrel{(*)}{=} \\ &= c_1 q_1 + \max_{w_1 \in \{\underline{w}_1, \overline{w}_1\}} \big(h_1 + J_2^\star\big)(x_1 + q_1 + w_1) \end{aligned} \tag{7}$$

where, in step $(*)$, we use the definition (4b) and the fact that the maximum of a convex function over the compact set

[2] Such constraints may lead to infeasible problems. For example: $T = 1, x_1 = 0, u_1 \in [0,1], w_1 \in [0,1], x_2 \in [5,10]$.

[3] $\tilde{h}_k$ are convex since $\mathbf{1}([L_{k+1}, U_{k+1}])$, the indicator function of a convex set, is convex [20], and the sum of convex functions remains convex.

$[\underline{w}_1, \overline{w}_1]$ is always reached at the boundaries of the set. Next, we introduce the affine cost $z_1(w_1) \stackrel{\text{def}}{=} z_{1,0} + z_{1,1}\, w_1$, where we constrain $z_{1,i}$ to satisfy the two linear equations

$$z_{1,0} + z_{1,1}\, w = h_1(x_1 + q_1 + w), \quad \forall w \in \{\underline{w}_1, \overline{w}_1\}.$$

Note that, for fixed x_1 and q_1, the function $z_1(w_1)$ is a linear interpolation of the mapping $w_1 \mapsto h_1(x_1 + q_1 + w_1)$, matching the value at points $\{\underline{w}_1, \overline{w}_1\}$. Since h_1 is convex, the linear interpolation dominates it, so condition (6b) is satisfied. Furthermore, by (7), J_{mM} is achieved for $w_1 \in \{\underline{w}_1, \overline{w}_1\}$, so condition (6c) is also obeyed.

Having checked the induction at time $k = 1$, we now assume that the statements of Theorem 3.1 are true for times $t = 1, \ldots, k$. Equation (6c) written for stage k then yields:

$$J_{mM} = \max_{(\theta_1, \theta_2) \in \Theta} \left[\theta_1 + J^{\star}_{k+1}(\theta_2) \right], \tag{8a}$$

$$\Theta \stackrel{\text{def}}{=} \Big\{ (\theta_1, \theta_2) : \theta_1 \stackrel{\text{def}}{=} \sum_{t=1}^{k} \big(c_t \cdot q_t(\boldsymbol{w}_t) + z_t(\boldsymbol{w}_{t+1}) \big), \quad \theta_2 \stackrel{\text{def}}{=} x_1 + \sum_{t=1}^{k} \big(q_t(\boldsymbol{w}_t) + w_t \big) \Big\}. \tag{8b}$$

Since $q_t(\cdot)$ and $z_t(\cdot)$ are affine functions, this implies that, although the uncertainties $w_1, \ldots, w_k$ lie in a set with 2^k vertices, i.e., the hyper-rectangle $\mathscr{H}_{k+1} = \mathscr{W}_1 \times \cdots \times \mathscr{W}_k$, they are only able to affect the objective J_{mM} through (two) affine combinations, taking values in the set Θ. Such a polyhedron, arising as a 2-dimensional affine projection of a k-dimensional hyper-rectangle, is called a *zonogon*. It belongs to a larger class of polytopes, known a *zonotopes*, whose combinatorial structure and properties are well documented in the discrete and computational geometry literature (see Chapter 7 of [21] for an excellent introduction).

The main properties of a zonogon that we are interested in are summarized in Lemma 1.1, found in the Appendix. In particular, Θ is centrally symmetric, and has at most $2k$ vertices (see example in Fig. 1). Furthermore, by numbering the vertices of Θ in counter-clockwise fashion, starting at

$$\boldsymbol{v}_0 \equiv \boldsymbol{v}^- \stackrel{\text{def}}{=} \arg\max_{\theta_1} \{ \arg\min_{\theta_2} \{ \boldsymbol{\theta} \in \Theta \} \}, \tag{9}$$

we can establish the following result concerning the points of Θ that are relevant in our problem:

Lemma 4.1: The maximum value in optimization problem (8a) is achieved for $(\theta_1, \theta_2) \in \{\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_k\}$.

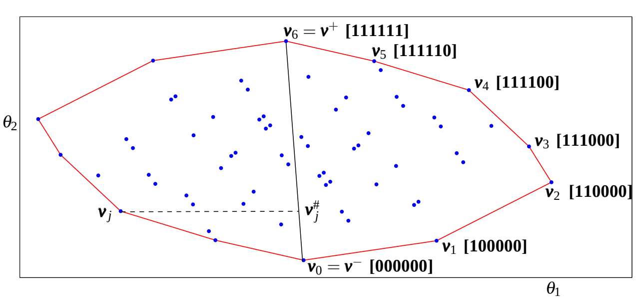

Fig. 1: Zonogon obtained from projecting a hypercube in $\mathbb{R}^6$.

Proof: By convexity, the maximum in (8a) is reached at $\boldsymbol{v} \in \text{ext}(\Theta)$. Furthermore, any extreme point on the left, $\boldsymbol{v}_j$, is dominated by some point to its right, $\boldsymbol{v}_j^{\#} \in \text{conv}(\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k)$, having $\theta_1 [\boldsymbol{v}_j^{\#}] \geq \theta_1 [\boldsymbol{v}_j]$ (see Fig. 1 for an illustration). ■

Since the argument presented in the lemma will be recurring throughout several of our proofs and constructions, we end this subsection by introducing two useful definitions, and generalizing the previous result.

Consider the system of coordinates (θ_1, θ_2) in $\mathbb{R}^2$, and let $\mathscr{S} \subset \mathbb{R}^2$ denote an arbitrary, finite set of points and $\mathscr{P}$ denote any (possibly non-convex) polygon such that its set of vertices is exactly $\mathscr{S}$. With $\boldsymbol{y}^- \stackrel{\text{def}}{=} \arg\max_{\theta_1} \{ \arg\min_{\theta_2} \{ \boldsymbol{\theta} \in \mathscr{S} \} \}$, and $\boldsymbol{y}^+ \stackrel{\text{def}}{=} \arg\max_{\theta_1} \{ \arg\max_{\theta_2} \{ \boldsymbol{\theta} \in \mathscr{S} \} \}$, by numbering the vertices of the convex hull of $\mathscr{S}$ in a counter-clockwise fashion, starting at $\boldsymbol{y}_0 \stackrel{\text{def}}{=} \boldsymbol{y}^-$, and with $\boldsymbol{y}_m = \boldsymbol{y}^+$, we define the *right side* of $\mathscr{P}$ and the *zonogon hull* of $\mathscr{S}$ as follows:

Definition 1: The *right side* of an arbitrary polygon $\mathscr{P}$ is:

$$\text{r-side}(\mathscr{P}) \stackrel{\text{def}}{=} \{ \boldsymbol{y}_0, \boldsymbol{y}_1, \ldots, \boldsymbol{y}_m \}.$$

Definition 2: The *zonogon hull* of a set of points $\mathscr{S}$ is:

$$\text{z-hull}(\mathscr{S}) \stackrel{\text{def}}{=} \Big\{ \boldsymbol{y} \in \mathbb{R}^2 : \boldsymbol{y} = \boldsymbol{y}_0 + \sum_{i=1}^{m} w_i \cdot (\boldsymbol{y}_i - \boldsymbol{y}_{i-1}),\, w_i \in [0, 1] \Big\}.$$

As the name suggests, r-side$(\mathscr{P})$ represents the vertices on the right side of $\mathscr{P}$. The zonogon hull of a set $\mathscr{S}$ is a zonogon having exactly the same vertices on the right side as the convex hull of $\mathscr{S}$, i.e., $\text{r-side}\big(\text{z-hull}(\mathscr{S})\big) = \text{r-side}\big(\text{conv}(\mathscr{S})\big)$. Some examples of zonogon hulls are shown in Fig. 2 (note that the initial points in $\mathscr{S}$ do not necessarily fall inside the zonogon hull, and, as such, there is no general inclusion relation between the zonogon hull and the convex hull). The reason for introducing this object is that it allows for the following generalization of Lemma 4.1:

Corollary 4.2: If $\mathscr{P}$ is any polygon in $\mathbb{R}^2$ with a finite set $\mathscr{S}$ of vertices, and $g : \mathbb{R} \to \mathbb{R}$ is any convex function, then the following equalities hold for $f(\boldsymbol{\theta}) \stackrel{\text{def}}{=} \theta_1 + g(\theta_2)$:

$$\max_{\boldsymbol{\theta} \in \mathscr{P}} f(\boldsymbol{\theta}) = \max_{\boldsymbol{\theta} \in \text{conv}(\mathscr{P})} f(\boldsymbol{\theta}) = \max_{\boldsymbol{\theta} \in \text{r-side}(\mathscr{P})} f(\boldsymbol{\theta}) =$$

$$\max_{\boldsymbol{\theta} \in \mathscr{S}} f(\boldsymbol{\theta}) = \max_{\boldsymbol{\theta} \in \text{z-hull}(\mathscr{S})} f(\boldsymbol{\theta}) = \max_{\boldsymbol{\theta} \in \text{r-side}(\text{z-hull}(\mathscr{S}))} f(\boldsymbol{\theta}).$$

Proof: The proof is identical to that of Lemma 4.1, and is omitted for brevity. ■

Using this result, whenever we are faced with a maximization of a convex function $\theta_1 + g(\theta_2)$, we can switch between different feasible sets, without affecting the overall optimal value of the problem.

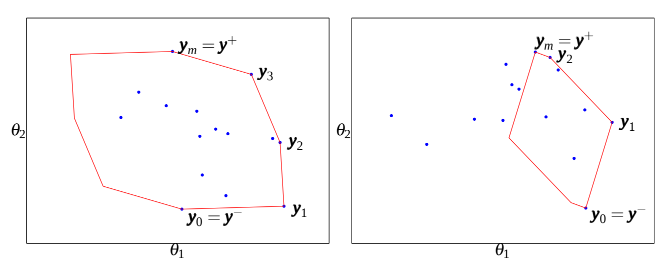

Fig. 2: Examples of zonogon hulls for different sets $\mathscr{S} \in \mathbb{R}^2$.

In the context of Lemma 4.1, the above result allows us to restrict attention from a potentially large set of relevant points (the 2^k vertices of the hyper-rectangle $\mathscr{H}_{k+1}$), to the $k+1$ vertices found on the right side of the zonogon Θ. This also gives insight into why computing an affine controller $q_{k+1}(\cdot)$ with $k+1$ degrees of freedom, yielding the same objective function value J_{mM}, might actually be possible.

In the remaining part of Section IV-A, we would like to further narrow down this set of relevant points, by using the structure and properties of the optimal control law $u^\star_{k+1}(x_{k+1})$ and optimal value function $J^\star_{k+1}(x_{k+1})$, derived in Section II. Before proceeding, we first reduce the notational clutter by introducing several simplifications and assumptions.

1) Simplified Notation and Assumptions: To start, we omit the time subscripts (k or $k+1$) whenever possible, and take:

$$\begin{aligned} \theta_1(\boldsymbol{w}) &\stackrel{\text{def}}{=} a_0 + \boldsymbol{a}'\boldsymbol{w}, \quad \theta_2(\boldsymbol{w}) \stackrel{\text{def}}{=} b_0 + \boldsymbol{b}'\boldsymbol{w}, \\ q_{k+1}(\boldsymbol{w}) &\equiv q(\boldsymbol{w}) \stackrel{\text{def}}{=} q_0 + \boldsymbol{q}'\boldsymbol{w}, \end{aligned} \tag{10}$$

where $\boldsymbol{a} = (a_1,\dots,a_k)$ and $\boldsymbol{b} = (b_1,\dots,b_k)$ are the *generators* of the zonogon Θ. We use the same counter-clockwise numbering of the vertices of Θ as introduced in Section IV-A, i.e., $\boldsymbol{v}_0 \equiv \boldsymbol{v}^-,\dots,\boldsymbol{v}_p \equiv \boldsymbol{v}^+,\dots,\boldsymbol{v}_{2p} = \boldsymbol{v}^-$, where $2p$ is the number of vertices of Θ.

Also, since $\theta_2 \equiv x_{k+1}$, instead of referring to $J^\star_{k+1}(x_{k+1})$ and $u^\star_{k+1}(x_{k+1})$, we use $J^\star(\theta_2)$ and $u^\star(\theta_2)$, and we use the short-hand notations $u^\star(\boldsymbol{v_i})$, $J^\star(\boldsymbol{v_i})$ and $g(\boldsymbol{v}_i)$, instead of $u^\star(\theta_2[\boldsymbol{v}_i])$, $J^\star(\theta_2[\boldsymbol{v}_i])$ and $g(\theta_2[\boldsymbol{v_i}] + u^\star(\theta_2[\boldsymbol{v}_i]))$, respectively. We also make the following simplifying assumptions:

Assumption 1: The uncertainty vector at time $k+1$ belongs to the unit hypercube of $\mathbb{R}^k$, i.e., $\mathscr{H}_{k+1} = [0,1]^k$.

Assumption 2: The zonogon Θ has a maximal number of vertices, i.e., $p = k$.

Assumption 3: The vertex of the hypercube projecting to $\boldsymbol{v}_i$, $i \in \{0,\dots,k\}$, is exactly $[1,1,\dots,1,0,\dots,0]$, i.e., 1 in the first k components and 0 thereafter (see Fig. 1).

These assumptions are made only to facilitate the exposition, and result in no loss of generality. To see this, note that Assumption 1 can always be achieved by adequate translation and scaling of the generators $\boldsymbol{a}$ and $\boldsymbol{b}$, and Assumption 3 can be satisfied by renumbering the coordinates of the hyper-rectangle. As for Assumption 2, we argue that an extension of our construction to the degenerate case $p < k$ is immediate (one could also remove the degeneracy by applying an infinitesimal perturbation to the generators $\boldsymbol{a}$ or $\boldsymbol{b}$, with infinitesimal cost implications).

2) Further Analysis of the Induction Hypothesis: With the new notation, substituting $J^\star(\theta_2) = c\cdot u^\star(\theta_2) + g(\theta_2 + u^\star(\theta_2))$ in (8a) yields the following optimization problem:

$$(OPT) \quad \begin{aligned} J_{mM} &= \max_{(\gamma_1^\star,\gamma_2^\star)\in\Gamma^\star} \Big[\ \gamma_1^\star + g(\gamma_2^\star)\ \Big], \\ \Gamma^\star &\stackrel{\text{def}}{=} \Big\{ (\gamma_1^\star,\gamma_2^\star) : \gamma_1^\star \stackrel{\text{def}}{=} \theta_1 + c\cdot u^\star(\theta_2), \\ &\qquad \gamma_2^\star \stackrel{\text{def}}{=} \theta_2 + u^\star(\theta_2),\ (\theta_1,\theta_2)\in\Theta \Big\}. \end{aligned} \tag{11}$$

A characterization for the set $\Gamma^\star = (\gamma_1^\star,\gamma_2^\star)$ can be obtained by replacing the optimal control law[4] from (5a):

$$(\gamma_1^\star,\gamma_2^\star) = \begin{cases} (\theta_1 + c\cdot U,\ \theta_2 + U), & \text{if } \theta_2 < y^\star - U \\ (\theta_1 - c\cdot\theta_2 + c\cdot y^\star,\ y^\star), & \text{otherwise} \\ (\theta_1 + c\cdot L,\ \theta_2 + L), & \text{if } \theta_2 > y^\star - L. \end{cases} \tag{12}$$

In particular, note that the optimal control law discriminates points $\boldsymbol{\theta} = (\theta_1,\theta_2) \in \Theta$ based on their position relative to the horizontal band:

$$\mathscr{B}_{LU} \stackrel{\text{def}}{=} \big\{\ (\theta_1,\theta_2)\in\mathbb{R}^2\ :\ \theta_2 \in [y^\star - U, y^\star - L]\ \big\}. \tag{13}$$

Since $\theta_2 \equiv x_{k+1}$, $\mathscr{B}_{LU}$ exactly corresponds to the state-space region when the bound constraints on the control at time $k+1$ are non-binding, while points with $\theta_2 < y^\star - U$ or $\theta_2 > y^\star - L$ fall in regions where the upper or lower bounds are binding, respectively.

The set $\Gamma^\star$ is, in general, not necessarily polyhedral (see Fig. 3 for an example). However, the following compact characterization for the maximizers in Problem (OPT) from (11) is immediate.

Lemma 4.3: The maximum in Problem (OPT) over $\Gamma^\star$ is reached on the right side of $\Delta_{\Gamma^\star} \stackrel{\text{def}}{=} \operatorname{conv}\big(\{\boldsymbol{y}_0^\star,\dots,\boldsymbol{y}_k^\star\}\big)$, where $\boldsymbol{y}_i^\star$, $i\in\{0,\dots,k\}$, are given by:

$$\boldsymbol{y}_i^\star \stackrel{\text{def}}{=} \big(\gamma_1^\star(\boldsymbol{v}_i),\gamma_2^\star(\boldsymbol{v}_i)\big) = \big(\theta_1[\boldsymbol{v}_i] + cu^\star(\boldsymbol{v}_i), \theta_2[\boldsymbol{v}_i] + u^\star(\boldsymbol{v}_i)\big). \tag{14}$$

Proof: By Lemma 4.1, the maximum in (8a) is reached at one of the vertices $\boldsymbol{v}_0,\boldsymbol{v}_1,\dots,\boldsymbol{v}_k$ of the zonogon Θ. Since this problem is equivalent to Problem (OPT) in (11), we can immediately conclude that the maximum of the latter optimization is reached at the points $\{\boldsymbol{y}_i^\star\}_{1\le i\le k}$ given by (14). Furthermore, since $g(\cdot)$ is convex (see Property **P2** of the optimal DP solution, in Section II), we can apply Corollary 4.2, and replace the points $\boldsymbol{y}_i^\star$ with the right side of their convex hull, r-side$(\Delta_{\Gamma^\star})$, without changing the result of the optimization problem, which completes the proof. ■

B. Construction of the Affine Control Law.

Having analyzed the consequences of the induction hypothesis in Theorem 3.1, we now return to the task of completing the inductive proof. In the current section, we exhibit an affine control law $q_{k+1}(\boldsymbol{w}_{k+1})$ that is robustly feasible, and that leaves the overall min-max cost J_{mM} unchanged, when used at time $k+1$ in conjunction with the original convex state cost, $h_{k+1}(x_{k+2})$.

In the simplified notation, the problem to be solved is to find an affine control law $q(\boldsymbol{w})$ such that:

$$\begin{aligned} J_{mM} &= \max_{\boldsymbol{w}\in\mathscr{H}_{k+1}} \Big[\ \theta_1 + c\cdot q(\boldsymbol{w}) + g\big(\theta_2 + q(\boldsymbol{w})\big)\ \Big] \\ L &\le q(\boldsymbol{w}) \le U\ , \quad \forall\, \boldsymbol{w}\in\mathscr{H}_{k+1}. \end{aligned}$$

The maximization represents the problem solved by the disturbances, when the affine controller, $q(\boldsymbol{w})$, is used instead of the optimal controller, $u^\star(\theta_2)$. The first equation

[4]For simplicity, we focus on the case when $g(\cdot)$ has a unique minimizer, such that $\underline{y} = \overline{y} = y^\star$ in (5a), (5b).

amounts to ensuring that the overall objective, J_{mM}, remains unchanged, and the inequalities are a restatement of the robust feasibility condition. The system can be rewritten:

$$(AFF)\quad \begin{aligned} J_{mM} &= \max_{(\gamma_1,\gamma_2)\in\Gamma} \left[\ \gamma_1 + g(\gamma_2)\ \right] \\ L &\le q(\boldsymbol{w}) \le U\ , \end{aligned} \tag{15}$$

$$\text{where}\quad \Gamma \stackrel{\text{def}}{=} \{(\gamma_1,\gamma_2) : \gamma_1 \stackrel{\text{def}}{=} \theta_1 + c\cdot q(\boldsymbol{w}), \\ \gamma_2 \stackrel{\text{def}}{=} \theta_2 + q(\boldsymbol{w}), (\theta_1,\theta_2)\in\Theta\} \tag{16}$$

Since $\theta_{1,2}$ and q are all affine in $\boldsymbol{w}$, by (10), the set Γ represents a new zonogon in $\mathbb{R}^2$, with generators $\boldsymbol{a} + c\cdot\boldsymbol{q}$ and $\boldsymbol{b}+\boldsymbol{q}$. Furthermore, since g is convex, Problem (AFF) over Γ is analogous to that in (8a), defined over the zonogon Θ. Thus, in analogy with the discussion in Section IV-A (Lemma 4.1), we conclude that the maximum in (AFF) occurs at a vertex of Γ found in r-side(Γ).

In a different sense, (AFF) is also similar to (OPT) in (11), in that the same convex objective, $\xi_1 + g(\xi_2)$, is maximized in both problems, but over different feasible sets, $\Gamma^\star$ for (OPT) and Γ for (AFF), respectively. From Lemma 4.3, the maximum in (OPT) is reached on r-side$(\Delta_{\Gamma^\star})$. From the discussion in the previous paragraph, the maximum in (AFF) occurs on r-side(Γ). Therefore, in order to compare the two problems, we must relate the sets r-side$(\Delta_{\Gamma^\star})$ and r-side(Γ).

In this context, we introduce the central idea behind the construction of the affine control law. Recalling the concept of a *zonogon hull* introduced in Definition 2, we argue that, if the affine coefficients of the controller, $q_0, \boldsymbol{q}$, were computed in such a way that the zonogon Γ actually corresponded to the *zonogon hull* of the set $\{\boldsymbol{y}_0^\star, \boldsymbol{y}_1^\star, \ldots, \boldsymbol{y}_k^\star\}$, then, by Corollary 4.2, we could immediately conclude that the optimal values in (OPT) and (AFF) are the same. This motivates the following algorithm.

Algorithm 1: Compute affine controller $q(\boldsymbol{w})$

Require: $\theta_1(\boldsymbol{w}), \theta_2(\boldsymbol{w}), g(\cdot), u^\star(\cdot)$

1: **if** (Θ falls below $\mathscr{B}_{LU}$) **or** ($\Theta \subseteq \mathscr{B}_{LU}$) **or** ($\Theta$ falls above $\mathscr{B}_{LU}$) **then**
2: Return $q(\boldsymbol{w}) = u^\star(\theta_2(\boldsymbol{w}))$.
3: **else**
4: Apply mapping (12) to obtain $\boldsymbol{y}_i^\star, i \in \{0,\ldots,k\}$.
5: Compute the set $\Delta_{\Gamma^\star} = \text{conv}\left(\{\boldsymbol{y}_0^\star,\ldots,\boldsymbol{y}_k^\star\}\right)$.
6: Let r-side$(\Delta_{\Gamma^\star}) = \{\boldsymbol{y}_0^\star, \boldsymbol{y}_1^\star, \ldots, \boldsymbol{y}_s^\star\} \cup \{\boldsymbol{y}_t^\star\} \cup \{\boldsymbol{y}_r^\star, \ldots, \boldsymbol{y}_k^\star\}$ be the set of points on the right side of $\Delta_{\Gamma^\star}$.
7: Solve the following system for $q_0,\ldots,q_k$ and K_U, K_L:

$$\begin{cases} q_0 + \cdots + q_i = u^\star(\boldsymbol{v}_i), & \forall\, \boldsymbol{y}_i^\star \in \text{r-side}(\Delta_{\Gamma^\star}) \\ \dfrac{a_i + c\cdot q_i}{b_i + q_i} = K_U, & i = s+1,\ldots,\min(t,r) \\ \dfrac{a_i + c\cdot q_i}{b_i + q_i} = K_L, & i = \max(t,s)+1,\ldots,r \end{cases} \tag{17}$$

8: Return $q(\boldsymbol{w}) = q_0 + \sum_{i=1}^k q_i w_i$.
9: **end if**

Proving that the algorithm produces the expected result is slightly technical - for complete details, we refer the interested reader to [17]. Here, we give intuition for the constraints in system (17), and the reasons why the construction works. In order to have the zonogon Γ be the same as the zonogon hull of $\{\boldsymbol{y}_0^\star,\ldots,\boldsymbol{y}_k^\star\}$, we must ensure that the vertices on the right side of Γ exactly correspond to the points on the right side of $\Delta_{\Gamma^\star} = \text{conv}\left(\{\boldsymbol{y}_0^\star,\ldots,\boldsymbol{y}_k^\star\}\right)$. This is achieved in two stages. First, we ensure that vertices $\boldsymbol{w}_i$ of the hypercube $\mathscr{H}_{k+1}$ that are mapped by the optimal control law $u^\star(\cdot)$ into points $\boldsymbol{y}_i^\star \in \text{r-side}(\Delta_{\Gamma^\star})$ (through the succession of mappings $\boldsymbol{w}_i \stackrel{(8b)}{\mapsto} \boldsymbol{v}_i \in \text{r-side}(\Theta) \stackrel{(14)}{\mapsto} \boldsymbol{y}_i^\star \in \text{r-side}(\Delta_{\Gamma^\star})$), will be mapped by the affine control law, $q(\cdot)$, into the same point $\boldsymbol{y}_i^\star$ (through the mappings $\boldsymbol{w}_i \stackrel{(8b)}{\mapsto} \boldsymbol{v}_i \in \text{r-side}(\Theta) \stackrel{(16)}{\mapsto} \boldsymbol{y}_i^\star \in \text{r-side}(\Delta_{\Gamma^\star})$). This is done in the first set of constraints, by *matching* the value of the optimal control law at any such points. Second, we ensure that any matched points $\boldsymbol{y}_i^\star$ actually correspond to the vertices on the right side of the zonogon Γ. This is done in the second and third set of constraints in (17), by computing the affine coefficients q_j in such a way that the resulting segments in the generators of the zonogon Γ, namely $\binom{a_j + c\cdot q_j}{b_j + q_j}$, are all *aligned*, i.e., form the same angle with the γ_1 axis, with cotangents given by the K_U, K_L variables, respectively. Geometrically, this exactly corresponds to the situation shown in Fig. 3.

We remark that the above algorithm does not explicitly require the control $q(\boldsymbol{w})$ to be robustly feasible, i.e., the second condition in (15). However, this property turns out to hold as a direct result of the way matching and alignment are performed in Algorithm 1.

C. Construction of the Affine State Cost

Note that we have essentially sketched only the first part of the induction step. For the second part, we need to show how an affine stage cost can be computed, such that constraints (6b) and (6c) are satisfied.

Returning temporarily to the notation with time indices, we remark that, in solving problem (AFF) of (15), we have shown that there exists an affine $q_{k+1}(\boldsymbol{w}_{k+1})$ such that:

$$\begin{aligned} J_{mM} &= \max_{\boldsymbol{w}_{k+1}\in\mathscr{H}_{k+1}} \Big[\gamma_1(\boldsymbol{w}_{k+1}) + g_{k+1}\big(\gamma_2(\boldsymbol{w}_{k+1})\big)\Big] \\ &\stackrel{(4b)}{=} \max_{\boldsymbol{w}_{k+2}\in\mathscr{H}_{k+2}} \Big[\pi_1(\boldsymbol{w}_{k+2}) + h_{k+2}\big(\pi_2(\boldsymbol{w}_{k+2})\big) + \\ &\qquad J^\star_{k+2}\big(\pi_2(\boldsymbol{w}_{k+2})\big)\Big], \end{aligned} \tag{18}$$

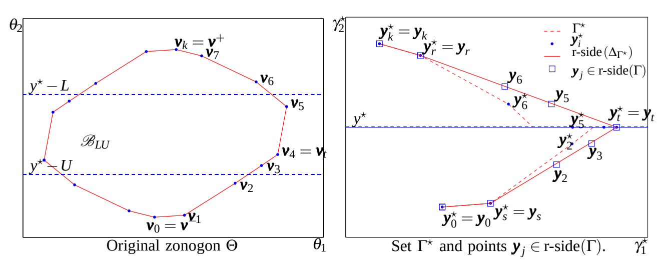

Fig. 3: Matching and alignment performed in Algorithm 1.

where $\pi_1(\boldsymbol{w}_{k+2}) \stackrel{\text{def}}{=} \gamma_1(\boldsymbol{w}_{k+1})$, and $\pi_2(\boldsymbol{w}_{k+2}) \stackrel{\text{def}}{=} \gamma_2(\boldsymbol{w}_{k+1}) + w_{k+2}$. Is is easy to note that

$$\Pi^\star \stackrel{\text{def}}{=} \big(\pi_1(\boldsymbol{w}_{k+2}), \pi_2(\boldsymbol{w}_{k+2})\big) \tag{19}$$

represents yet another zonogon, obtained by projecting the hyper-rectangle $\mathscr{H}_{k+2}$ into $\mathbb{R}^2$. In this context, the problem that remains to be solved is replacing the convex function $h_{k+2}\big(\pi_2(\boldsymbol{w}_{k+2})\big)$ with an affine function $z_{k+2}(\boldsymbol{w}_{k+2})$, such that the analogues of (6b) and (6c) are obeyed:

$$\begin{aligned} & z_{k+2}(\boldsymbol{w}_{k+2}) \ge h_{k+2}\big(\pi_2(\boldsymbol{w}_{k+2})\big), \qquad \forall\, \boldsymbol{w}_{k+2} \in \mathscr{H}_{k+2} \\ & J_{mM} = \max_{\boldsymbol{w}_{k+2} \in \mathscr{H}_{k+2}} \Big[\, \pi_1(\boldsymbol{w}_{k+2}) + z_{k+2}(\boldsymbol{w}_{k+2}) + J^\star_{k+2}\big(\pi_2(\boldsymbol{w}_{k+2})\big) \Big]. \end{aligned} \tag{20}$$

With the same simplified notation and analysis as presented for the affine control, we can define $\text{r-side}\big(\Pi^\star\big) \stackrel{\text{def}}{=} \big\{\boldsymbol{v}_0, \boldsymbol{v}_1, \dots, \boldsymbol{v}_{k+1}\big\}$, and introduce the following algorithm for computing the affine cost $z(\boldsymbol{w})$:

Algorithm 2: Compute affine stage cost $z(\boldsymbol{w})$

Require: $\pi_1(\boldsymbol{w}), \pi_2(\boldsymbol{w}), h(\cdot), J^\star(\cdot)$.

1: Compute $\boldsymbol{y}^\star_i \stackrel{\text{def}}{=} \big(\pi_1[\boldsymbol{v}_i] + h(\boldsymbol{v}_i),\ \pi_2[\boldsymbol{v}_i]\big),\ i = 0, \dots, k+1$.

2: Compute the set $\Delta_{\Pi^\star} = \text{conv}\big(\{\boldsymbol{y}^\star_0, \dots, \boldsymbol{y}^\star_{k+1}\}\big)$.

3: Let $\text{r-side}(\Delta_{\Pi^\star}) \stackrel{\text{def}}{=} \{\boldsymbol{y}^\star_{s(1)}, \dots, \boldsymbol{y}^\star_{s(n)}\}$, where $s(1) \le s(2) \le \dots \le s(n) \in \{0, \dots, k+1\}$ are sorted indices.

4: Solve the following system for z_j, $(j \in \{0, \dots, k+1\})$, and $K_{s(i)}$, $(i \in \{2, \dots, n\})$:

$$\begin{cases} \displaystyle\sum_{j=0}^{s(i)} z_j = h\left(\boldsymbol{v}_{s(i)}\right), & \forall\, \boldsymbol{y}^\star_{s(i)} \in \text{r-side}(\Delta_{\Pi^\star}) \\ \dfrac{z_j + a_j}{b_j} = K_{s(i)}, & \forall\, j \in \{s(i-1)+1, \dots, s(i)\}, \\ & \forall\, i \in \{2, \dots, n\}. \end{cases} \tag{21}$$

5: Return $z(\boldsymbol{w}) = z_0 + \sum_{i=1}^{k+1} z_i \cdot w_i$.

The key idea behind Algorithm 2 is exactly the same as that in the construction of the affine control law. In particular, with the coefficients of the affine cost, z_i, computed as above, the set $\big\{\big(\pi_1(\boldsymbol{w}) + z(\boldsymbol{w}), \pi_2(\boldsymbol{w})\big),\ \boldsymbol{w} \in \mathscr{H}_{k+2}\big\}$, characterizing the maximization problem (20) exactly represents the *zonogon hull* of the set of points $\big\{\boldsymbol{y}^\star_0, \dots, \boldsymbol{y}^\star_{k+1}\big\}$, and hence, by Corollary 4.2, the optimal value in (20) is the same as the optimal value in (18), namely J_{mM}. To visualize how the algorithm is working, an example is included in Fig. 4 below. Full details and a complete proof can be found in [17].

D. Counterexamples for potential extensions.

A natural question to ask is whether results similar to Theorem 3.1 could be extended to more general problems, such as multi-dimensional linear systems. In particular, whether affine policies in the disturbances remain optimal, and whether affine relaxations of the stage costs are possible, without loss of optimality.

Unfortunately, it turns out that simple counterexamples can be found (see [17]) to answer both of the above questions negatively. To conclude, policies that are affine in the disturbances are, in general, *suboptimal* for problems with multiple dimensions, and replacing the convex state costs by (larger) affine costs would, in general, result in even *further* deterioration of the objective.

V. CONCLUSIONS AND FUTURE WORK

We have presented a novel approach for theoretically handling robust, multi-stage decision problems. The method utilized the connections between the geometrical properties of the feasible sets (zonogons), and the objective functions being optimized, so as to prune the set of relevant points and derive properties about the optimal policies for the problem.

One immediate direction of research would be to study systems with mixed constraints (on both state and control), and to explore extensions of the approach to nonconvex and nonlinear cost structures. Another potential area of interest would be to quantify the performance of affine policies even in problems where they are known to be suboptimal. This could potentially lead to fast approximation algorithms, with solid theoretical foundations.

APPENDIX

Zonogons are all centrally symmetric, 2-dimensional $2p$-gons, arising as the projection of p-cubes to the plane. An example is shown in Fig. 1 of Section IV-A. These are the main objects of interest in our treatment, and the following lemma summarizes their most important properties:

Lemma 1.1: Let $\mathscr{H}_k = [0,1]^k$ be a k-dimensional hypercube, $k \ge 2$. For fixed $\boldsymbol{a}, \boldsymbol{b} \in \mathbb{R}^k$ and $a_0, b_0 \in \mathbb{R}$, consider the affine transformation $\pi : \mathbb{R}^k \to \mathbb{R}^2$, $\pi(\boldsymbol{w}) = \begin{bmatrix} \boldsymbol{a}' \\ \boldsymbol{b}' \end{bmatrix} \cdot \boldsymbol{w} + \begin{bmatrix} a_0 \\ b_0 \end{bmatrix}$ and the zonogon $\Theta \subset \mathbb{R}^2$ given by:

$$\Theta = \pi\left(\mathscr{H}_k\right) \stackrel{\text{def}}{=} \big\{\boldsymbol{\theta} \in \mathbb{R}^2 : \exists\, \boldsymbol{w} \in \mathscr{H}_k \text{ s.t. } \boldsymbol{\theta} = \pi(\boldsymbol{w})\big\}.$$

If we let $\mathscr{V}_\Theta$ denote the set of vertices of Θ, then the following properties are true:

1) $\exists\, \boldsymbol{O} \in \Theta$ such that Θ is symmetric around $\boldsymbol{O}$: $\forall\, \boldsymbol{x} \in \Theta \Rightarrow 2\boldsymbol{O} - \boldsymbol{x} \in \Theta$.
2) $|\mathscr{V}_\Theta| = 2p \le 2k$ vertices. Also, $p < k$ if and only if $\exists\, i \ne j \in \{1, \dots, k\}$ such that $\text{rank}\left(\begin{bmatrix} a_i & a_j \\ b_i & b_j \end{bmatrix}\right) < 2$.
3) If we number the vertices of $\mathscr{V}_\Theta$ in cyclic order:

$$\mathscr{V}_\Theta = \left(\boldsymbol{v}_0, \dots, \boldsymbol{v}_i, \boldsymbol{v}_{i+1}, \dots, \boldsymbol{v}_{2p-1}\right), \quad \left(\boldsymbol{v}_{2p+i} \stackrel{\text{def}}{=} \boldsymbol{v}_{(2p+i) \mod (2p)}\right)$$

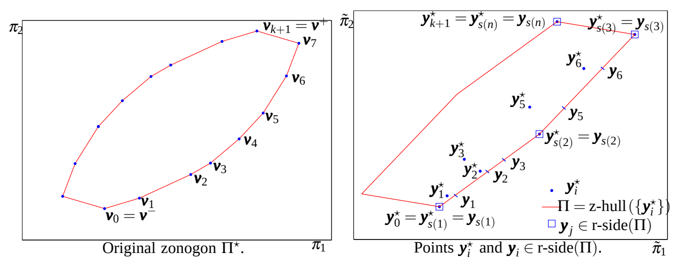

Fig. 4: Matching and alignment performed in Algorithm 2.

then $2\boldsymbol{O} - \boldsymbol{v}_i = \boldsymbol{v}_{i+p}$, and we have the following representation for Θ as a Minkowski sum of line segments:

$$\Theta \stackrel{\text{def}}{=} \boldsymbol{O} + \sum_{i=1}^{p} \lambda_i \cdot \frac{\boldsymbol{v}_i - \boldsymbol{v}_{i-1}}{2}, \quad -1 \leq \lambda_i \leq 1.$$

4) If $\exists\, \boldsymbol{w}_{1,2} \in \mathscr{H}_k$ such that $\boldsymbol{v}_1 \stackrel{\text{def}}{=} \pi(\boldsymbol{w}_1) = \boldsymbol{v}_2 \stackrel{\text{def}}{=} \pi(\boldsymbol{w}_2)$ and $\boldsymbol{v}_{1,2} \in \mathscr{V}_\Theta$, then $\exists\, j \in \{1,\ldots,k\}$ such that $a_j = b_j = 0$.
5) With the same numbering from (iii) and $k = p$, for any $i \in \{0,\ldots,2p-1\}$, the vertices of the hypercube that are projecting to $\boldsymbol{v}_i$ and $\boldsymbol{v}_{i+1}$, respectively, are adjacent, i.e. they only differ in exactly one component.

REFERENCES

[1] G. E. Dullerud and F. Paganini, *A Course in Robust Control Theory*. Springer, 2005.
[2] K. Zhou and J. C. Doyle, *Essentials of Robust Control*. Prentice Hall, 1998.
[3] A. Ben-Tal and A. Nemirovski, "Robust Convex Optimization," *Mathematics of Operations Research*, vol. 23, no. 4, pp. 769–805, 1998. [Online]. Available: http://mor.journal.informs.org/cgi/content/abstract/23/4/769
[4] L. El-Ghaoui, F. Oustry, and H. Lebret, "Robust Solutions to Uncertain Semidefinite Programs," *SIAM J. Optimization*, vol. 9, no. 1, pp. 33–52, 1998.
[5] D. Bertsimas and M. Sim, "Robust Discrete Optimization and Network Flows," *Mathematical Programming*, vol. 98, no. 1-3, pp. 49–71, 2003.
[6] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific, 2001.
[7] A. Bemporad, F. Borrelli, and M. Morari, "Min-Max Control of Constrained Uncertain Discrete-Time Linear Systems," *IEEE Transactions on Automatic Control*, vol. 48, no. 9, pp. 1600–1606, 2003.
[8] R. T. Rockafellar and R. J.-B. Wets, "The optimal recourse problem in discrete time: L1-multipliers for inequality constraints," *SIAM Journal of Control and Optimization*, vol. 16, pp. 16–36, 1978.
[9] J. Löfberg, "Approximations of closed-loop minimax MPC," *Proceedings of the 42nd IEEE Conference on Decision and Control*, vol. 2, pp. 1438–1442 Vol.2, Dec. 2003.
[10] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski, "Adjustable robust solutions of uncertain linear programs," *Math. Program.*, vol. 99, no. 2, pp. 351–376, 2004.
[11] A. Ben-Tal, S. Boyd, and A. Nemirovski, "Control of Uncertainty-Affected Discrete Time Linear Systems via Convex Programming," 2005, working paper.
[12] E. Kerrigan and J. Maciejowski, "Properties of a New Parameterization for the Control of Constrained Systems with Disturbances," *Proceedings of the 2004 American Control Conference*, vol. 5, pp. 4669–4674 vol.5, June-July 2004. [Online]. Available: http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=01384049
[13] J. Skaf and S. Boyd, "Nonlinear Q-Design for Convex Stochastic Control," 2008, submitted to *IEEE Transactions on Automatic Control*. [Online]. Available: http://www.stanford.edu/~boyd/papers/nonlin_Q_param.html
[14] D. Bertsimas and D. Brown, "Constrained Stochastic LQC: A Tractable Approach," *IEEE Transactions on Automatic Control*, vol. 52, no. 10, pp. 1826–1841, Oct. 2007.
[15] P. Grieder, P. A. Parrilo, and M. Morari, "Robust Receding Horizon Control - Analysis & Synthesis," in *Proceedings of the 42nd IEEE Conference on Decision and Control*, vol. 1, Dec. 2003, pp. 941–946 Vol.1.
[16] P. J. Goulart and E. Kerrigan, "Relationships Between Affine Feedback Policies for Robust Control with Constraints," in *Proceedings of the 16th IFAC World Congress on Automatic Control*, July 2005.
[17] D. Bertsimas, D. Iancu, and P. Parrilo, "Optimality of Affine Policies in Multi-stage Robust Optimization," 2009, submitted for publication. [Online]. Available: http://arxiv.org/abs/0904.3986
[18] A. Ben-Tal, B. Golany, A. Nemirovski, and J.-P. Vial, "Retailer-Supplier Flexible Commitments Contracts: A Robust Optimization Approach," *Manufacturing & Service Operations Management*, vol. 7, no. 3, pp. 248–271, 2005.
[19] A. Ben-Tal and A. Nemirovski, "Robust optimization - methodology and applications," *Math. Program.*, vol. 92, no. 3, pp. 453–480, 2002.
[20] T. Rockafellar, *Convex Analysis*. Princeton University Press, 1970.
[21] G. Ziegler, *Lectures on Polytopes*. Springer, 2nd edition, 2003.