

On the Adversarial Convex Body Chasing Problem

Yue Guan Longxu Pan Daigo Shishika Panagiotis Tsiotras

Abstract—In this work, we extend the convex body chasing problem to an adversarial setting, where an agent (the Player) is tasked to chase a sequence of convex bodies generated adversarially by an opponent. The Player aims to minimize the cost associated with its total movements, while the Opponent tries to maximize. The set of feasible convex bodies is finite and known to both agents, which allows us to provide performance guarantees with max-min optimality rather than via the competitive ratio. Under some mild assumptions, we show the continuity of the optimal value function and provide an algorithm to numerically generate the optimal policies with performance guarantees. Finally, the theoretical results are verified through numerical examples.

I. INTRODUCTION

The Convex Body Chasing (CBC) problem was proposed in [1] to study the interaction between convexity and metrical task systems. Later, it was realized that many problems can be viewed as variants of the CBC problem, including scheduling [2], efficient covering [3], safe machine-learned advice [4], self-organizing lists [5], and the k-server problem [6]. In the CBC problem, an *online* agent receives a request sequence of T convex sets $Q_1, \dots, Q_T$ in a normed space $\mathcal{X}$ of dimension d . The agent starts at $\mathbf{x}_0$, and at time step t observes the set Q_t and then moves to a new point $\mathbf{x}_t \in Q_t$, paying a cost $\|\mathbf{x}_t - \mathbf{x}_{t-1}\|$. The objective of the agent is to minimize the total cost at the final time T . Previous works [7]–[9] aimed to design online algorithms that ensure that the total cost exceeds the best hindsight cost by at most a factor α , known as the “competitive ratio.” The most recent work [9] has achieved a competitive ratio of order $\mathcal{O}(\sqrt{d \log T})$, despite the fact that the proposed online algorithm chooses $\mathbf{x}_t$ without knowledge of the future sets $Q_{t+1}, \dots, Q_T$.

In the classic CBC problem, the agent needs to select a point that balances the performance for all possible future convex sets. Consequently, the *competitive ratio* is considered as the performance metric for most of the existing CBC solution algorithms. In many real-world scenarios, however, the convex bodies are selected (potentially adversarially) from a known collection of convex sets. Specifically, the recently introduced dynamic Defender-Attacker Blotto game [10] possesses such structure. With this additional




Fig. 1. Timeline of the aCBC game. At each time step t , the Opponent first selects a state i_t to assign a corresponding convex set $Q_t^{(i_t)}$ for the Player. The Player then selects a point $\mathbf{x}_t$ within the assigned convex set.

information, one expects to obtain better performance guarantees than competitive ratios.

In this work, we propose a two-player adversarial convex bodies chasing (aCBC) game, where the Player is tasked to chase the sequence of convex bodies generated by the Opponent. Different from the original formulation, we assume that the bodies are selected from a finite collection of convex regions, which is given to both the Player and the Opponent at the beginning of the game. The adversarial selection of the convex bodies is further constrained over a graph, which implies that the currently selected convex body has an impact on the convex bodies available at the next time step¹. The Player’s movement is also constrained within its (compact) reachable set and is dependent on its current position². We formulate this adversarial problem as a zero-sum sequential game [11], where the Player aims to minimize its total cost (associated with its total movement), while the Opponent tries to maximize it.

The contribution of this work is threefold: (i) the novel formulation of the adversarial CBC game; (ii) derivation of the Lipschitz continuity of the max-min value function under certain assumptions; and (iii) a proposed numerical algorithm that provides ϵ -suboptimal performance guarantees with respect to the optimal max-min solution.

II. PROBLEM FORMULATION

The adversarial convex body chasing (aCBC) game is played sequentially between two agents: the Player and the Opponent. The game evolves over a compact subset $\mathcal{X}$ of a normed Euclidean space $\mathbb{R}^d$. At each time step, a convex region $Q_t \subseteq \mathcal{X}$ is selected by the Opponent, and then the Player must choose a point $\mathbf{x}_t \in Q_t$. The Player tries to minimize its total cost $\sum_{t=0}^{T-1} c(\mathbf{x}_t, \mathbf{x}_{t+1})$ over a finite horizon T for some non-negative continuous cost function c , while the Opponent aims to maximize this total cost. The sequence of decisions for the two agents of the game is presented in Fig. 1.

We use $\mathbf{x}_t$ to denote the state of the Player at time t and treat $\mathcal{X}$ as the state space of the Player. Different from the classical CBC problem, we restrict the Player’s selection of its next state to a neighborhood of its current

Yue Guan is a PhD Student with the School of Aerospace Engineering, Georgia Institute of Technology. Email: yguan44@gatech.edu

Longxu Pan is an MS student with the School of Aerospace Engineering, Georgia Institute of Technology. Email: lpan64@gatech.edu

Daigo Shishika is an Assistant Professor with the Department of Mechanical Engineering Department, Georgia Mason University. Email: dshishik@gmu.edu

Panagiotis Tsiotras is the David & Andrew Lewis Chair and Professor with the School of Aerospace Engineering, Georgia Institute of Technology. Email: tsiotras@gatech.edu

*The first two authors contributed equally.

¹One can remove the graph constraint by using a fully connected graph.

²Since there are only finitely many compact convex bodies, one can remove this constraint by setting the reachable set to be a convex hull that contains all convex bodies.

state. Specifically, we require that $\mathbf{x}_{t+1} \in \mathcal{R}(\mathbf{x}_t)$ for all $t = 0, \dots, T-1$, where $\mathcal{R}(\mathbf{x}_t)$ denotes the *reachable set* of the Player from its current state $\mathbf{x}_t$. We use $\mathcal{R} : \mathcal{X} \mapsto 2^{\mathcal{X}}$ to denote the reachability correspondence (a set-valued map), where $2^{\mathcal{X}}$ denotes the power set of $\mathcal{X}$. Correspondences will be denoted with the symbol $\rightsquigarrow$. In this case, we will write $\mathcal{R} : \mathcal{X} \rightsquigarrow \mathcal{X}$.

Assumption 1. For all $\mathbf{x} \in \mathcal{X}$, $\mathcal{R}(\mathbf{x})$ is solid, compact and convex.

At time t , we assign a state $i_t \in \mathcal{V}$ to the Opponent, which evolves over a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Given the Opponent's current state i_t , the Opponent can move to any one of the neighboring nodes i_{t+1} , such that $(i_t, i_{t+1}) \in \mathcal{E}$. We denote the set of all neighbors of node i_t as $\mathcal{N}_{i_t}$. Before the game starts, we assign the Opponent a finite collection of convex regions $\mathcal{Q} = \{\mathcal{Q}_t^{(i)}\}_{t=0, i=1}^{T, |\mathcal{V}|}$. At each time step t , the Opponent selects a convex region $\mathcal{Q}_t^{(i_t)}$ from $\mathcal{Q}$ by selecting its state i_t . Once a convex region $\mathcal{Q}_t^{(i_t)}$ is selected at time t by the Opponent, the Player needs to move to a point $\mathbf{x}_t \in \mathcal{R}(\mathbf{x}_{t-1}) \cap \mathcal{Q}_t^{(i_t)}$. For ease of notation, we use $\Theta_t^{(i_t)}(\mathbf{x}_{t-1})$ to denote the intersection $\mathcal{R}(\mathbf{x}_{t-1}) \cap \mathcal{Q}_t^{(i_t)}$. We make the following two assumptions on the collection $\mathcal{Q}$ and the information structure of the game.

Assumption 2. For all $i \in \mathcal{V}$ and $t = 0, \dots, T$, the set $\mathcal{Q}_t^{(i)} \subseteq \mathcal{X}$ is solid, compact and convex.

Assumption 3. The collection $\mathcal{Q}$ and the graph $\mathcal{G}$ are common knowledge to both agents.

An illustrative example of the proposed aCBC game is presented in Fig. 2.

With these assumptions, the major difference of the aCBC from the original CBC formulation is that (i) the feasible convex bodies are compact, and (ii) the set of feasible convex bodies in aCBC game is *finite* and is *common knowledge* to both agents, (iii) the convex body selection is restricted by the Opponent's dynamics on a graph. The compactness assumption and the finiteness of the feasible set of convex bodies significantly reduce the Opponent's freedom of selecting convex regions. This allows the construction of a min-max solution that is discussed in the next sections.

To ensure the existence of a min-max solution, we further make the following assumptions on the continuity of the cost function c and the reachability correspondence $\mathcal{R}$.

Assumption 4. The cost $c : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$ is continuous.

To characterize the continuity of correspondences, we first assign a metric to the power set $2^{\mathcal{X}}$.

Definition 1 (Hausdorff distance [12]). Given two subsets E and F of the normed space $\mathcal{X}$, the Hausdorff distance between E and F is defined as

$$\text{dist}_H(E, F) = \max \left\{ \sup_{\mathbf{x} \in E} \inf_{\mathbf{y} \in F} \|\mathbf{x} - \mathbf{y}\|, \sup_{\mathbf{y} \in F} \inf_{\mathbf{x} \in E} \|\mathbf{x} - \mathbf{y}\| \right\}.$$

Assumption 5. The reachability correspondence $\mathcal{R} : \mathcal{X} \rightsquigarrow \mathcal{X}$ is continuous under the Hausdorff distance.

To avoid degeneracy, we make the following assumption, so that the aCBC game is an optimization problem rather than a feasibility problem.

Assumption 6. For all time steps $t = 0, \dots, T-1$, the collection of sets $\mathcal{Q}$ satisfies:

$$\text{int}(\mathcal{R}(\mathbf{x}_t) \cap \mathcal{Q}_{t+1}^{(i_{t+1})}) \neq \emptyset, \mathbf{x}_t \in \mathcal{Q}_t^{(i_t)}, i_{t+1} \in \mathcal{N}_{i_t}. \quad (1)$$

We consider Markov policies for both agents. The Player's (deterministic) policy at time $t \geq 1$ is given by $\pi_t(\mathbf{x}_{t-1}, i_t) \in \mathcal{R}(\mathbf{x}_{t-1}) \cap \mathcal{Q}_t^{(i_t)}$, which explicitly considers the convex body $\mathcal{Q}_t^{(i_t)}$ selected by the Opponent at time t and the Player's reachability constraint. The Opponent's policy at time $t \geq 1$ is given by $\sigma_t(\mathbf{x}_{t-1}, i_{t-1}) \in \mathcal{N}_{i_{t-1}}$, which reflects the graph constraint. At time $t = 0$, the Opponent is allowed to freely select a node $i_0 \in \mathcal{V}$ according to its policy $\sigma_0(\mathcal{Q}, \mathcal{G})$. Then, the Player selects a point $\mathbf{x}_0 \in \mathcal{Q}_0^{(i_0)}$ according to $\pi_0(i_0)$, after observing the Opponent's selection i_0 . We denote the strategies used by the Player and the Opponent as $\pi = \{\pi_t\}_{t=0}^T$ and $\sigma = \{\sigma_t\}_{t=0}^T$ respectively.

Problem 1. Given a graph $\mathcal{G}$, a collection of compact convex bodies $\mathcal{Q} = \{\mathcal{Q}_t^{(i)}\}_{t=0, i=1}^{T, |\mathcal{V}|}$, a cost function c , and a reachability correspondence $\mathcal{R}$, find the optimal total cost of the aCBC game C^* along with the corresponding optimal strategies for the Player and the Opponent under the information structure in Assumption 3.

To solve this Problem, we follow a value-based approach, where, at each decision point, the agent computes its optimal "cost-to-go." The rest of the paper will address the technical details regarding the solution of this optimization problem.

III. VALUE FUNCTIONS

We introduce two optimal value functions: $V_t(\mathbf{x}_{t-1}, i_t)$ for the Player and $U_t(\mathbf{x}_{t-1}, i_{t-1})$ for the Opponent. The two optimal value functions will be computed through a backward induction scheme [13].

A. Backward Induction

At the terminal time step T , the Player has knowledge of its previous state $\mathbf{x}_{T-1}$ and the Opponent state i_T , and the Player is about to make its final move. Since there are no moves after time T , the Player only needs to consider the immediate cost. Consequently, we have

$$V_T(\mathbf{x}_{T-1}, i_T) = \min_{\mathbf{x}_T \in \Theta_T^{(i_T)}(\mathbf{x}_{T-1})} c(\mathbf{x}_{T-1}, \mathbf{x}_T). \quad (2)$$

For time steps $t \in \{T-1, \dots, 1\}$, the Player needs to optimize its selection of a new state $\mathbf{x}_t$ to minimize both the immediate cost and the future cost. The Player needs to further consider the fact that the Opponent will observe its move and then best-respond with i_{t+1} to maximize the future cumulative costs. Consequently, the optimal value for the Player is computed as

$$V_t(\mathbf{x}_{t-1}, i_t) = \min_{\mathbf{x}_t \in \Theta_t^{(i_t)}(\mathbf{x}_{t-1})} \left\{ c(\mathbf{x}_{t-1}, \mathbf{x}_t) + \max_{i_{t+1} \in \mathcal{N}_{i_t}} V_{t+1}(\mathbf{x}_t, i_{t+1}) \right\}. \quad (3)$$




Fig. 2. Example of an aCBC game with a 4-node graph and two-dimensional Player state $\mathcal{X}$. At time step 0, the Opponent starts at Node 2 and the Player selects a point $\mathbf{x}_0$ within $Q_0^{(2)}$. The Opponent then moves to Node 4 and the Player selects the point $\mathbf{x}_1 \in \mathcal{R}(\mathbf{x}_0) \cap Q_1^{(4)}$, inducing a cost of $c(\mathbf{x}_0, \mathbf{x}_1)$. The Opponent then moves to Node 3 and the Player moves to $\mathbf{x}_2 \in \mathcal{R}(\mathbf{x}_1) \cap Q_2^{(3)}$ incurring a cost of $c(\mathbf{x}_1, \mathbf{x}_2)$, and the game continues.

Finally, for the initial Player state selection at $t = 0$, there is no reachability constraint. As a result, the value function only depends on the initial Opponent state i_0 , and we have

$$V_0(i_0) = \min_{\mathbf{x}_0 \in Q_0^{(i_0)}} \max_{i_1 \in \mathcal{N}_{i_0}} V_1(\mathbf{x}_0, i_1). \quad (4)$$

The Opponent's value function is constructed similarly to that of the Player and is given by

$$U_T(\mathbf{x}_{T-1}, i_{T-1}) = \max_{i_T \in \mathcal{N}_{i_{T-1}}} \min_{\mathbf{x}_T \in \Theta_T^{(i_T)}(\mathbf{x}_{T-1})} c(\mathbf{x}_{T-1}, \mathbf{x}_T), \quad (5)$$

$$U_t(\mathbf{x}_{t-1}, i_{t-1}) = \max_{i_t \in \mathcal{N}_{i_{t-1}}} \min_{\mathbf{x}_t \in \Theta_t^{(i_t)}(\mathbf{x}_{t-1})} \left\{ c(\mathbf{x}_{t-1}, \mathbf{x}_t) + U_{t+1}(\mathbf{x}_t, i_t) \right\}, \quad (6)$$

$$U_0(\mathcal{G}, \mathcal{Q}) = \max_{i_0 \in \mathcal{V}} \min_{\mathbf{x}_0 \in Q_0^{(i_0)}} U_1(\mathbf{x}_0, i_0). \quad (7)$$

Remark 1. For the aCBC game, the Opponent's value U_0 is equal to the optimal game value C^* .

B. Continuity of the Value Functions

The first question one may ask regarding the value functions is whether the use of min operator in the value propagation rules is justified. To answer this question, we first show that the intersection correspondence Θ_t^i is continuous, and then we prove the continuity of the value functions.

Assumption 5 only regards the continuity of the reachability correspondence $\mathcal{R}$ rather than of the intersection correspondence Θ_t^i . To bridge this gap, we present the following lemma to guarantee the continuity of Θ_t^i .

Lemma 1. Let $\Gamma : \mathcal{X} \rightsquigarrow \mathcal{Y}$ be continuous, and let $\Gamma(\mathbf{x})$ be compact and convex for all $\mathbf{x} \in \mathcal{X}$. Consider a closed convex set $F \subseteq \mathcal{Y}$ such that $\text{int}(\Gamma(\mathbf{x}) \cap F) \neq \emptyset$ for all $\mathbf{x} \in C$, where $C \subseteq \mathcal{X}$ is closed. Then, the correspondence $\Xi : C \rightsquigarrow \mathcal{Y}$ defined by $\Xi(\mathbf{x}) = \Gamma(\mathbf{x}) \cap F$ is continuous.

Proof. The proof is given in the Appendix. $\square$

The following lemma provides some insight into the continuity of marginal functions of the following form:

$$\phi(\mathbf{x}) = \min_{\mathbf{y} \in \Gamma(\mathbf{x})} f(\mathbf{x}, \mathbf{y}). \quad (8)$$

Lemma 2 (Proposition 2.9 in [12]). Consider a continuous function $f : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$ and a continuous correspondence $\Gamma : \mathcal{X} \rightsquigarrow \mathcal{X}$. If Γ is compact-valued, then the marginal function $\phi : \mathcal{X} \mapsto \mathbb{R}$ in (8) is continuous.

Note that the value functions in (3) and (6) take the form of a marginal function. Consequently, we can prove the following theorem on the continuity of the value functions using Lemma 2.

Theorem 3. For all $t \in \{1, \dots, T\}$, $i_{t-1} \in \mathcal{V}$ and $i_t \in \mathcal{N}_{i_{t-1}}$, the optimal value functions $V_t(\mathbf{x}_{t-1}, i_t)$ and $U_t(\mathbf{x}_{t-1}, i_{t-1})$ are continuous with respect to their first argument on $Q_{t-1}^{(i_{t-1})}$.

Proof. The proof is based on an inductive argument. We only present the key idea for the induction step here. See the full version [14] for details.

Assume that the two value functions $V_{t+1}(\mathbf{x}_t, i_{t+1})$ and $U_{t+1}(\mathbf{x}_t, i_t)$ are continuous with respect to their first argument on $Q_t^{(i_t)}$ for all $i_t, i_{t-1} \in \mathcal{V}$ such that $i_t \in \mathcal{N}_{i_{t-1}}$. From the assumption, the function $g(\mathbf{x}_t) = \max_{i_{t+1} \in \mathcal{N}_{i_t}} V_{t+1}(\mathbf{x}_t, i_{t+1})$ is continuous. Together with the continuity of the cost c , the function $f(\mathbf{x}_{t-1}, \mathbf{x}_t) = c(\mathbf{x}_{t-1}, \mathbf{x}_t) + g(\mathbf{x}_t)$ is also continuous. Since the intersection $\Theta_t^{(i_t)}(\mathbf{x}_{t-1})$ is compact, the value function in (3) can be written as $V_t(\mathbf{x}_{t-1}, i_t) = \min_{\mathbf{x}_t \in \Theta_t^{(i_t)}(\mathbf{x}_{t-1})} f(\mathbf{x}_{t-1}, \mathbf{x}_t)$, and the minimization is well-defined. From Lemma 1, we have that the intersection correspondence $\Theta_t^{(i_t)}$ is continuous. The continuity of the Player value function then follows directly from Lemma 2. A similar analysis applies to the Opponent value function. $\square$

C. Optimal Policies

With the optimal value functions computed, the optimal policy of the Player and the Opponent can be obtained via

$$\begin{aligned} \pi_t^*(\mathbf{x}_{t-1}, i_t) &\in \underset{\mathbf{x}_t \in \Theta_t^{(i_t)}(\mathbf{x}_{t-1})}{\operatorname{argmin}} \left\{ c(\mathbf{x}_{t-1}, \mathbf{x}_t) + U_{t+1}(\mathbf{x}_t, i_t) \right\}, \\ \sigma_T^*(\mathbf{x}_{t-1}, i_{t-1}) &\in \underset{i_t \in \mathcal{N}_{i_{t-1}}}{\operatorname{argmax}} V_t(\mathbf{x}_{t-1}, i_t). \end{aligned}$$

IV. ALGORITHMIC SOLUTION

To numerically compute the optimal values, we propose an algorithm that meshes the domain $\mathcal{X}$ and approximate U_t and V_t at the vertices of the mesh. In order to derive approximation error bounds induced by the discretization, we first strengthen the continuity results on U_t and V_t .

A. Lipschitz Continuity of the Value Functions

The following two assumptions are imposed to ensure the Lipschitz continuity of the value functions.

Assumption 7. The cost function c is L_c -Lipschitz continuous. That is, for all $(\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}') \in \mathcal{X} \times \mathcal{X}$, the following holds

$$|c(\mathbf{x}, \mathbf{y}) - c(\mathbf{x}', \mathbf{y}')| \leq L_c(\|\mathbf{x} - \mathbf{x}'\| + \|\mathbf{y} - \mathbf{y}'\|).$$

Assumption 8. The correspondence $\Theta_t^{(i)}$ is L_Θ -Lipschitz continuous. That is, for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ and $i \in \mathcal{V}$, the following holds

$$\text{dist}_H(\Theta_t^{(i)}(\mathbf{x}), \Theta_t^{(i)}(\mathbf{x}')) \leq L_\Theta \|\mathbf{x} - \mathbf{x}'\|.$$

Theorem 4. Under Assumptions 7 and 8, it follows that, for all $t \in \{1, \dots, T\}$, $i_{t-1} \in \mathcal{V}$, and $i_t \in \mathcal{N}_{i_{t-1}}$, the optimal value functions $V_t(\mathbf{x}_{t-1}, i_t)$ and $U_t(\mathbf{x}_{t-1}, i_{t-1})$ are both $L_{v,t}$ -Lipschitz, where the Lipschitz constant is given by

$$L_{v,t} = L_c \sum_{k=1}^{T-t+1} (1 + L_\Theta)^k. \quad (9)$$

Proof. One can first extend the continuity result in Lemma 2 to Lipschitz continuity with the strengthened assumptions. The proof of Theorem 4 can then be constructed using a similar inductive argument presented in the proof of Theorem 3. For the detailed proof, please see [14]. $\square$

B. Discretization Scheme and Discretized Value Functions

Theorem 4 implies that the Lipschitz constant $L_{v,t}$ decreases monotonically as the time step t approaches the horizon T . Naturally, finer discretization resolution is preferred at the beginning of the game to ensure low approximation error. Consequently, we allow different resolutions at different time steps. We use $\delta_{\mathcal{X},t}$ to denote the discretization size of the domain $\mathcal{X}$ at time step t , and we denote the set of vertices on the mesh at time t by $\hat{\mathcal{X}}_t = \{\hat{\mathbf{x}}_t^k\}_{k=1}^{M_t}$. To ensure that the domain is properly discretized with a discretization size $\delta_{\mathcal{X},t}$, we require

$$\min_{\hat{\mathbf{x}}_t \in \hat{\mathcal{X}}_t} \|\hat{\mathbf{x}}_t - \mathbf{x}_t\| \leq \delta_{\mathcal{X},t}, \quad \mathbf{x}_t \in \mathcal{X}, \quad t = 0, \dots, T; \quad (10)$$

$$\min_{\hat{\mathbf{x}}_0 \in \hat{\mathcal{X}}_0 \cap \mathcal{Q}_0^{(i_0)}} \|\hat{\mathbf{x}}_0 - \mathbf{x}_0\| \leq \delta_{\mathcal{X},0}, \quad \mathbf{x}_0 \in \mathcal{Q}_0^{(i_0)}, i_0 \in \mathcal{V}; \quad (11)$$

$$\min_{\hat{\mathbf{x}}_t \in \hat{\mathcal{X}}_t \cap \mathcal{R}(\hat{\mathbf{x}}_{t-1}) \cap \mathcal{Q}_t^{(i_t)}} \|\hat{\mathbf{x}}_t - \mathbf{x}_t\| \leq \delta_{\mathcal{X},t}, \quad (12)$$

$$\mathbf{x}_t \in \mathcal{R}(\hat{\mathbf{x}}_{t-1}) \cap \mathcal{Q}_t^{(i_t)}, \quad t = 1, \dots, T, \\ \hat{\mathbf{x}}_{t-1} \in \hat{\mathcal{X}}_{t-1} \cap \mathcal{Q}_{t-1}^{(i_{t-1})}, i_t \in \mathcal{N}_{i_{t-1}}.$$

In the above criteria, we first require that the mesh has the required resolution over the whole domain $\mathcal{X}$ as

in (10). We further require that the resolution within the intersections of the reachable sets and the convex bodies is fine enough as in (11) and (12). An example of the discretization is presented in Fig. 3. The black vertices are constructed to discretize the domain $\mathcal{X}$ according to (10). For the mesh consisting only black vertices, there is no vertex in $\mathcal{Q}_0^{(1)}$. Consequently, two blue vertices are added to $\hat{\mathcal{X}}_0$ to satisfy (11). Furthermore, although the intersection $\mathcal{R}(\hat{\mathbf{x}}_0) \cap \mathcal{Q}_1^{(2)}$ contains one black vertex, the mesh does not have a fine enough resolution within the intersection. As a result, two red vertices are added to satisfy (12).




Fig. 3. An example of the discretization scheme.

With the discretization of $\mathcal{X}$, the Player restricts its action selection $\mathbf{x}_t$ at time step t to the vertices $\hat{\mathcal{X}}_t$, which leads to a natural decrease in its performance. We consider the worst-case scenario, where the Opponent knows the discretization mesh used by the Player. The resultant value functions then take the vertices of the mesh $\hat{\mathcal{X}}_t$ as the discretized state space for the Player. We denote the discretized value functions as $\hat{V}$ and $\hat{U}$ for the Player and the Opponent, respectively. For example, the Player's value function at time t is given by

$$\hat{V}_t(\hat{\mathbf{x}}_{t-1}, i_t) = \min_{\hat{\mathbf{x}}_t \in \hat{\mathcal{X}}_t \cap \Theta_t^{(i_t)}(\hat{\mathbf{x}}_{t-1})} \left\{ c(\hat{\mathbf{x}}_{t-1}, \hat{\mathbf{x}}_t) + \max_{i_{t+1} \in \mathcal{N}_{i_t}} \hat{V}_{t+1}(\hat{\mathbf{x}}_t, i_{t+1}) \right\}, \quad (13)$$

where $\hat{\mathbf{x}}_{t-1} \in \hat{\mathcal{X}}_{t-1}$. For the detailed definitions of the discretized value functions, please see [14].

We denote the optimal strategies induced from the discretized value functions as $\hat{\pi}^*$ and $\hat{\sigma}^*$ for the Player and the Opponent, respectively. For example, the Player's discretized optimal policy at time step $1 \leq t \leq T-1$ is defined as

$$\hat{\pi}_t^*(\hat{\mathbf{x}}_{t-1}, i_t) \in \underset{\hat{\mathbf{x}}_t \in \hat{\mathcal{X}}_t \cap \Theta_t^{(i_t)}(\hat{\mathbf{x}}_{t-1})}{\text{argmin}} \{c(\hat{\mathbf{x}}_{t-1}, \hat{\mathbf{x}}_t) + \hat{U}_{t+1}(\hat{\mathbf{x}}_t, i_t)\}.$$

and similarly for the policy of the Opponent.

C. Error Bounds of Discretization

In this section we provide bounds for the discretization error $\varepsilon = |\hat{U}_0(\mathcal{G}, \mathcal{Q}) - U_0(\mathcal{G}, \mathcal{Q})|$. Once the relation between the discretization size and the error is established, we will present an algorithm that solves aCBC games numerically.

Based on the Lipschitz continuity properties of the value functions, one expects to obtain a better approximation of the optimal value function with a finer mesh. Since the error due to discretization propagates over time, it is relatively unclear how fine a discretization is needed at each time step to provide the desired performance bound. We answer the above question in the following theorem.

Theorem 5. Given a discretization scheme $\{\delta_{\mathcal{X},t}\}_{t=0}^T$ that satisfies (10) – (12), the difference between the discretized value function and the optimal value function is bounded, for all $t \in \{1, \dots, T\}$ as follows:

$$U_t(\hat{\mathbf{x}}_{t-1}, i_{t-1}) \leq \hat{U}_t(\hat{\mathbf{x}}_{t-1}, i_{t-1}) \leq U_t(\hat{\mathbf{x}}_{t-1}, i_{t-1}) + \hat{\epsilon}_t, \quad (14a)$$

$$V_t(\hat{\mathbf{x}}_{t-1}, i_t) \leq \hat{V}_t(\hat{\mathbf{x}}_{t-1}, i_t) \leq V_t(\hat{\mathbf{x}}_{t-1}, i_t) + \hat{\epsilon}_t, \quad (14b)$$

where, $\hat{\epsilon}_t = \sum_{\tau=t}^{T-1} (L_c + L_{v,\tau+1})\delta_{\mathcal{X},\tau} + L_c\delta_{\mathcal{X},T}$.

Proof. We only present a sketch of the induction step for the proof (see [14] for details). Assume that (14) holds at $t+1$. Fix a vertex $\hat{\mathbf{x}}_{t-1} \in \hat{\mathcal{X}}_{t-1}$. From the definition of $\hat{V}_t$,

$$\begin{aligned} \hat{V}_t(\hat{\mathbf{x}}_{t-1}, i_t) &= \min_{\hat{\mathbf{x}}_t \in \hat{\mathcal{X}}_t \cap \Theta_t^{(i_t)}(\hat{\mathbf{x}}_{t-1})} \{c(\hat{\mathbf{x}}_{t-1}, \hat{\mathbf{x}}_t) + \hat{U}_{t+1}(\hat{\mathbf{x}}_t, i_t)\} \\ &\leq \min_{\hat{\mathbf{x}}_t \in \hat{\mathcal{X}}_t \cap \Theta_t^{(i_t)}(\hat{\mathbf{x}}_{t-1})} \{c(\hat{\mathbf{x}}_{t-1}, \hat{\mathbf{x}}_t) + U_{t+1}(\hat{\mathbf{x}}_t, i_t)\} + \hat{\epsilon}_{t+1} \\ &\leq \min_{\mathbf{x}_t \in \Theta_t^{(i_t)}(\hat{\mathbf{x}}_{t-1})} \{c(\hat{\mathbf{x}}_{t-1}, \mathbf{x}_t) + U_{t+1}(\mathbf{x}_t, i_t)\} + \hat{\epsilon}_{t+1} \\ &\quad + L_c\delta_{\mathcal{X},t} + L_{v,t+1}\delta_{\mathcal{X},t} \\ &= V_t(\hat{\mathbf{x}}_{t-1}, i_t) + \hat{\epsilon}_t, \end{aligned}$$

where the first inequality is the result of the inductive hypothesis. The second inequality leverages (12), which ensures that for all $\mathbf{x}_t$ in the intersection, there is a vertex in the intersection that is $\delta_{\mathcal{X},t}$ -close. This fact, along with the Lipschitz continuity of the optimal value function U_t and the cost c , yields the desired result.

The proof for the other direction and for the Opponent's value can be constructed in a similar manner. $\square$

As a direct consequence of Theorem 5, the following corollary provides error bounds for the game value after discretization.

Corollary 6. The optimal game value $\hat{U}_0(\mathcal{G}, \mathcal{Q})$ due to discretization differs from the optimal max-min value $U_0(\mathcal{G}, \mathcal{Q})$ by at most

$$\varepsilon(\delta_{\mathcal{X}}) = \sum_{\tau=1}^{T-1} (L_c + L_{v,\tau+1})\delta_{\mathcal{X},\tau} + L_{v,1}\delta_{\mathcal{X},0} + L_c\delta_{\mathcal{X},T}.$$

If the suboptimality bound ε is given, an easy way to construct a proper discretization scheme is to use the following mesh sizes

$$\begin{aligned} \delta_{\mathcal{X},0} &= \frac{\varepsilon}{(T+1)L_{v,1}}, \quad \delta_{\mathcal{X},T} = \frac{\varepsilon}{(T+1)L_c}, \\ \delta_{\mathcal{X},t} &= \frac{\varepsilon}{(T+1)(L_c + L_{v,t+1})}, \quad t = 1, \dots, T-1. \end{aligned} \quad (15)$$

V. NUMERICAL SIMULATIONS

For the sake of simplicity and for illustrative purposes, we consider an aCBC game on a two-dimensional space where all sets are boxes. For some real numbers $a_1 < b_1$ and $a_2 < b_2$, we define the two-dimensional box set as the Cartesian product $[a_1, b_1] \times [a_2, b_2]$. The reachability correspondence is defined as $\mathcal{R}(\mathbf{x}) = \{\mathbf{y} \in \mathcal{X} \text{ s.t. } \|\mathbf{x} - \mathbf{y}\|_{\infty} \leq \rho\}$, for some $\rho > 0$. Taking Euclidean norm as the cost function, one can verify that the cost function c is 1-Lipschitz. Furthermore, for all time steps t and for all nodes i_t , the correspondence $\Theta_t^{(i_t)}$




Fig. 4. Illustrative aCBC instance with nested convex regions. The left figure the convex sets. The bottom right shows the optimal discretized trajectory. The top right presents the 2-node graph used in this scenario.

is 1-Lipschitz continuous under Hausdorff distance³. Given these two Lipschitz constants, one can use (15) to derive the desired discretization sizes $\delta_{\mathcal{X}}$ for a given suboptimality requirement ε .

We first verify that the discretized algorithm indeed converges to an optimal solution. The “nested” convex region example in Fig. 4 has a simple optimal solution, where the Player starts at any point within the intersection of all convex regions (marked in grey) and does not move. Under this optimal strategy, regardless of the actions of the Opponent, the Player can achieve zero cost.

However, as the discretization mesh $\hat{\mathcal{X}}_t$ changes over time, there may not be a point that is a vertex for all intermediate meshes. Consequently, the Player may move slightly from time to time under the computed optimal *discretized* policy, incurring a discretization error (see the zoom-in plot in Fig. 4). Since the optimal value U_0 is zero, the discretized value $\hat{U}_0$ is the discretization error. To verify that the discretization scheme in (15) indeed achieves the required performance, we run the algorithm with different suboptimality requirements ε and plotted the corresponding discretization errors in Fig. 5. One can see that the discretization error diminishes as ε decreases to zero. Furthermore, the discretization error is bounded by ε , which validates the bounds we derived in Corollary 6.

Next, we present a more complicated scenario with a graph of three nodes and a two-step horizon. Fig. 6 shows the




Fig. 5. Suboptimality requirement vs. discretization error plot.

³For two box sets E and Q , such that $E \cap Q \neq \emptyset$, one can show that $\min_{\mathbf{y} \in Q} \|\mathbf{x} - \mathbf{y}\|_2 = \min_{\mathbf{y} \in Q \cap E} \|\mathbf{x} - \mathbf{y}\|_2$ for all $\mathbf{x} \in E$. The rest of the analysis is then straightforward. See [14] for details.

convex regions $\mathcal{Q}$ and the graph $\mathcal{G}$. In this visualization, a convex region is color-filled if it is selected by the Opponent. The darker color marks the intersection of the reachable set and the corresponding convex set. For example, the darker region in $\mathcal{Q}_2^{(3)}$ (filled green) depicts $\mathcal{R}(\mathbf{x}_1) \cap \mathcal{Q}_2^{(3)}$, where $\mathbf{x}_1$ is the point selected in $\mathcal{Q}_1^{(1)}$ (filled blue).

The trajectory (in black color) in Fig. 6 is the max-min trajectory induced by the discretized optimal policies computed under using $\varepsilon = 0.2$. The trajectory of the Opponent is $1 \rightarrow 1 \rightarrow 3$. One can see that at the initial time step, the Player selects a vertical mid-point between $\mathcal{Q}_1^{(1)}$ and $\mathcal{Q}_1^{(3)}$ to balance the two possible convex regions at the next time step. As a result, the Opponent is indifferent between nodes 1 and 3 at $t = 1$.

Finally, Fig. 7 depicts the scenarios when one of the agents deviates from its discretized optimal strategy. One can see an increase in the cost when the Player unilaterally deviates and a decrease in the cost when the Opponent unilaterally deviates. Notice that when the Player deviates and does not select the mid-point in $\mathcal{Q}_1^{(1)}$, the Opponent selects $\mathcal{Q}_3^{(1)}$ to exploit the Player's error and maximize the cost.

VI. CONCLUSION

We have extended the convex body chasing problem to an adversarial setting, where a Player chases a sequence of convex bodies assigned adversarially by an Opponent. We showed that under the assumption that the set of convex bodies is finite and known to both agents, max-min optimal policies can be obtained. A discretization scheme was proposed to numerically solve for the max-min value function of the game, and performance guarantees were provided for the corresponding discretized policies. We will also leverage the aCBC framework to introduce cost to the dynamic Defender-Attacker Blotto Games [10].

APPENDIX

Proof of Lemma 1. From Corollary 2.12 in [12], we have that the intersection correspondence Ξ is upper semicontinuous (usc). It remains to show that Ξ is also lower




Fig. 6. The max-min trajectory from the discretized optimal policies.




Fig. 7. Trajectories with unilateral deviation. The left figure presents the case when the Player deviates and the right is when the Opponent deviates.

semicontinuous (lsc). To this end, let $\mathbf{x} \in C$, and consider an open set E such that $E \cap \Xi(\mathbf{x}) \neq \emptyset$. Since $\text{int}(\Xi(\mathbf{x}))$ is nonempty and convex, we have that $E \cap \text{int}(\Xi(\mathbf{x})) \neq \emptyset$ [15]. Let G be an open set such that

$$G \subseteq E \cap \text{int}(\Xi(\mathbf{x})) = E \cap \text{int}(F \cap \Gamma(\mathbf{x})). \quad (16)$$

Clearly, $G \cap \Gamma(\mathbf{x}) \neq \emptyset$. Consequently, from the lower semicontinuity of Γ , there exists an open neighborhood H around $\mathbf{x}$ such that for all $\mathbf{x}' \in H$, $G \cap \Gamma(\mathbf{x}') \neq \emptyset$. Since $G \subseteq F$ as in (16), we also have $G \cap \Gamma(\mathbf{x}') \cap F = G \cap \Xi(\mathbf{x}') \neq \emptyset$ for all $\mathbf{x}' \in H$. Finally, since $G \subseteq E$, it follows that, for all points $\mathbf{x}' \in H$, $E \cap \Xi(\mathbf{x}') \neq \emptyset$, which implies that Ξ is lower semicontinuous. Since Ξ is compact-valued, usc and lsc implies continuity under the Hausdorff distance (see Proposition 2.8 of [12]). $\square$

REFERENCES

- [1] J. Friedman and N. Linial, "On convex body chasing," *Discrete & Computational Geometry*, vol. 9, no. 3, pp. 293–321, 1993.
- [2] R. L. Graham, "Bounds for certain multiprocessing anomalies," *Bell System Technical Journal*, vol. 45, no. 9, pp. 1563–1581, 1966.
- [3] N. Alon, B. Awerbuch, and Y. Azar, "The online set cover problem," in *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, 2003, pp. 100–105.
- [4] A. Wei and F. Zhang, "Optimal robustness-consistency trade-offs for learning-augmented online algorithms," *Advances in Neural Information Processing Systems*, vol. 33, pp. 8042–8053, 2020.
- [5] D. D. Sleator and R. E. Tarjan, "Amortized efficiency of list update and paging rules," *Communications of the ACM*, vol. 28, no. 2, pp. 202–208, 1985.
- [6] E. Koutsoupias and C. H. Papadimitriou, "On the k-server conjecture," *Journal of the ACM*, vol. 42, no. 5, pp. 971–983, 1995.
- [7] S. Bubeck, Y. T. Lee, Y. Li, and M. Sellke, "Competitively chasing convex bodies," in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 2019, pp. 861–868.
- [8] S. Bubeck, B. Klartag, Y. T. Lee, Y. Li, and M. Sellke, "Chasing nested convex bodies nearly optimally," in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2020, pp. 1496–1508.
- [9] M. Sellke, "Chasing convex bodies optimally," in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2020, pp. 1509–1518.
- [10] D. Shishika, Y. Guan, M. Dorothy, and V. Kumar, "Dynamic defender-attacker Blotto game," *arXiv preprint arXiv:2112.09890*, 2021.
- [11] G. Owen, *Game Theory*. Emerald Group Publishing, 2013.
- [12] R. Freeman and P. V. Kokotović, *Robust Nonlinear Control Design: State-space and Lyapunov Techniques*. Springer Science & Business Media, 2008.
- [13] D. Bertsekas, *Dynamic Programming and Optimal Control: Volume I*. Athena scientific, 2012, vol. 1.
- [14] Y. Guan, L. Pan, D. Shishika, and P. Tsiotras, "Chasing convex bodies generated by an adversary," *arXiv preprint arXiv:2209.13606*, 2022.
- [15] T. Rockafellar, *Convex analysis*. Princeton university press, 1970, vol. 18.