Receding horizon games with stability guarantees

Emmanouil Dimou, Filippo Fabiani, and Alberto Bemporad

Abstract— **Receding-horizon games offer a very promising approach to extend nonzero-sum dynamic feedback games to handle local and shared constraints on the input variables in non-cooperative, multi-agent control systems. In this paper, we consider the multi-agent control of linear, discrete-time, time-invariant stabilizable systems subject to affine (shared) state and input constraints. We propose a game-theoretic model predictive control (MPC) formulation that makes the associated generalized Nash equilibrium (GNE) problem solvable at all time steps and the origin an exponentially stable equilibrium state of the overall closed-loop dynamics. While leveraging well-established results in solving jointly-convex linear-quadratic GNE problems, we exploit a polyhedral control Lyapunov function to enforce stability. The validity of the proposed receding horizon game scheme is demonstrated through simulations.**

I. Introduction

We address the problem of solving linear discrete-time dynamic games under linear input and state constraints in a receding horizon fashion. In this type of game models, multiple agents control their own subset of the input vector, in an optimal manner according to their own cost functions, while respecting local and shared constraints. In the so-called *receding-horizon game* (RHG), the agents at each discrete-time step measure the current state and, based on an available common model of the entire system dynamics, try to agree on some input strategies. The agreement means that every agent acts optimally over its own finite horizon criterion given the actions of all other agents, or, in other words, that the agents' strategies constitute a generalized Nash equilibrium (GNE). Receding-horizon approaches allow the agents to re-negotiate the decisions given the current state of the system, therefore achieving a feedback game-theoretic policy in spite of the fact that the agents solve an open-loop game at every time step, for numerical tractability reasons [1]. Indeed, receding horizon solutions have been proposed to dynamic games with application in autonomous driving [2], autonomous drone racing [3] and supply chains [4].

In RHGs, at every discrete time step the agents solve a GNE problem (GNEP) for the input strategies and apply only the first action of their own strategy to the system. Ensuring the solvability of the arising problems and convergence of the closed-loop dynamics to an equilibrium remains a fundamental challenge. In [5] the authors, under certain assumptions about the parameters of the underlying physical process, dealt with generalized potential games. Additionally, they provided sufficient conditions for the existence of a practically stable equilibrium point for the closed-loop system. In [6], the authors viewed the closed-loop system as a feedback interconnection of two dissipative systems, an open-loop strictly stable linear time invariant (LTI) system and a static nonlinearity, mapping the initial states of the system to the strategies in equilibrium. The stability of such interconnection was ensured by leveraging elements of monotone operator theory with the dissipativity property. In a similar setting, dissipativity-type conditions were also given in [7]. The authors in [1] observed that the strategies in equilibrium, in an open-loop information structure, of an infinite-horizon linear-quadratic (LQ) game coincide with the input sequences of linear-quadratic regulation (LQR) problems in a state space of higher dimension. This observation was based on a note, firstly introduced in [8]. They subsequently used the explicit expression of the infinite cost-to-go as a terminal cost in the receding horizon, constrained game and proved asymptotic convergence of the closed-loop system to an equilibrium. Game-theoretic MPC laws based on open-loop predictions under linear constraints and quadratic costs can be evaluated efficiently as LQ-GNEPs, see, e.g., [9]; explicit characterizations of the control law have been investigated in [10] by exploiting the properties of multi-parametric LQ games, with similar results obtained in [11].

It is also believed that RHGs are relevant to noncooperative, distributed MPC (DMPC) methods [1]. More specifically, in [12], the authors propose a "feasible cooperation" DMPC scheme in which the objectives of the agents, controlling the interconnected sub-systems are replaced by a common objective that guarantees some level of cooperation between the subsystems. At every discrete-time step, an algorithm solves interdependent quadratic programming problems iteratively, until it reaches a certain level of agreement between the subsystems and subsequently applies the initial control inputs. A similar approach was adopted by the authors in [13], who let the agents communicate several times while solving their local optimization problems. Furthermore, in [14] the authors proposed an a-posteriori verifiable stability condition, which was heavily relying on the degree of decoupling between the subsystems.

In this work, we consider game-theoretic MPC schemes for linear systems under linear input and state constraints based on open-loop predictions. Contrarily to a similar scheme proposed in [1], in which the infinite-horizon open-loop cost based on the agent's objectives is used as a terminal cost, we propose the use of a polyhedral control Lyapunov

This work was funded by the European Union (ERC Advanced Research Grant COMPACT, No. 101141351). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.
The authors are with the IMT School for Advanced Studies Lucca, Piazza San Francesco 19, 55100, Lucca, Italy (email: {emmanouil.dimou, filippo.fabiani, alberto.bemporad}@imtlucca.it).

function (PCLF) to ensure closed-loop stability. Imposing a PCLF for the closed-loop nicely maps into shared linear constraints, so that, when the agents' objectives are quadratic, we do not alter the LQ-nature of the resulting GNEP to solve at each time step.

The usage of PCLFs for the constrained stabilization of LTI systems traces back to [15], [16]. As clearly stated in [17], one of the reasons is that their domain of attraction (DoA) can approximate with arbitrary precision the maximal control positive-invariant (MCPI) set of the constrained LTI system. In linear MPC schemes, and more specifically in [18], the authors identify an infinity norm as a polyhedral Lyapunov function, while in [19] the authors apply a similar core idea in hybrid MPC schemes. Additionally, in [17] two distinct MPC methods with guarantees about the closed-loop stability of the origin were proposed, one using a PCLF, with a proper weight, as a terminal cost, and the other using the same PCLF to design input-affine contractive constraints. For a thorough survey on similar topics the interested reader may refer to [20].

The main contribution of this paper is to analyze the closed-loop properties of the resulting game-theoretic MPC scheme, namely the recursive existence of a GNE at every time step (and, therefore, the well-posedness of the control law) and the exponential convergence of the closed-loop dynamics to the origin from a specific set of initial conditions. The aforementioned additional shared constraints imposed in the game explicitly ask for a decrease in the value of a polyhedral function, which provably acts as a Lyapunov function for the closed-loop system on its DoA. Due to the affine nature of the proposed constraints, well-known results about the existence of GNE solutions [21] can be exploited.

This paper is organized as follows. In Section II we formalize the RHG the agents participate in. In Section III we review PCLFs and their corresponding DoA, which will serve as a basis for the stability analysis of our closed-loop system presented in Section IV. Section V shows the results of a numerical simulation of closed-loop game-theoretic MPC based on the proposed approach, while Section VI concludes the paper.

Notation: We denote by T the prediction horizon and by N the number of agents. The set of symmetric, positive definite (semi-definite), real matrices of dimensions $n \times n$ is denoted as $\mathbb{S}^n_{\succ(\succcurlyeq)0}$, the i-th row of a real matrix $A \in \mathbb{R}^{n\times m}$ as $A_i^\top$. Given N vectors $b_1, \dots, b_N \in \mathbb{R}^n$, $\mathrm{col}(b_1, \dots, b_N)$ denotes the stacked vector $\left[b_1^\top, \dots, b_N^\top\right]^\top$; given N real matrices $A_1, \cdots, A_N \in R^{n\times m}$ the block diagonal matrix with $A_1, \dots, A_N$ in the diagonal is denoted by $\mathrm{diag}(A_1, \dots, A_N)$. The subscript $_{\tau|k}$ indicates a τ-step ahead prediction of some quantity of interest made at time step k.

II. Problem formulation

We consider the following discrete-time stabilizable LTI system:

$$x_{k+1} = Ax_k + \sum_{i=1}^{N} B_i u_{k,i} \tag{1}$$

where $x_k \in \mathbb{R}^{n_x}$ is the state vector $u_{k,i} \in \mathbb{R}^{n_{u_i}}$ the control input commanded by agent i, $i = 1, \dots, N$. The dynamics are described by $A \in \mathbb{R}^{n_x \times n_x}$ and matrices $B_i \in \mathbb{R}^{n_x \times n_{u_i}}$, collectively stacked as $B = [B_1 \cdots B_N] \in \mathbb{R}^{n_x \times n_u}$, $n_u = \sum_{i=1}^{N} n_{u_i}$. The state variables and the control input of each agent are subject to the following polytopic constraints

$$x_k \in \mathcal{X} = \{x \in \mathbb{R}^{n_x} \mid A_x x \leq b_x\} \tag{2a}$$

$$u_{k,i} \in \mathcal{U}_i = \{u_i \in \mathbb{R}^{n_{u_i}} \mid A_{u,i} u_i \leq b_{u,i}\} \tag{2b}$$

with $A_x \in \mathbb{R}^{n_{x,c} \times n_x}$, $b_x \in \mathbb{R}^{n_{x,c}}$, $A_{u,i} \in \mathbb{R}^{n_{u,c_i} \times n_{u_i}}$, and $b_{u,i} \in \mathbb{R}^{n_{c_i}}$. Let $\mathcal{U} = \prod_{i\in\mathcal{I}} \mathcal{U}_i \subset \mathbb{R}^{n_u}$ be the polytope constraining the collective input $u_k = \mathrm{col}(u_{k,1}, \dots, u_{k,N}) \in \mathbb{R}^{n_u}$.

To pose the problem of interest, we make the following assumption on the system's dynamics and constraints:

Assumption 1: The pair (A, B) is stabilizable, $\mathcal{U}_i \neq \emptyset$ and $0 \in \mathcal{U}_i$, for all $i = 1, \dots, N$, $\mathcal{X} \neq \emptyset$, and $0 \in \mathcal{X}$.

Next, we formalize the non-cooperative dynamic Nash game the agents participate in. Specifically, we consider a RHG in which $\mathbf{u}_i = \mathrm{col}\left(u_{0|k,i}, \dots, u_{T-1|k,i}\right) \in \mathbb{R}^{Tn_{u_i}}$ denotes the sequence of control inputs commanded by agent i over a finite horizon of length $T \geq 1$, $\mathbf{u}$ is the collective control sequences of all agents, and $\mathbf{u}_{-i} = \mathrm{col}(\mathbf{u}_1, \dots, \mathbf{u}_{i-1}, \mathbf{u}_{i+1}, \dots, \mathbf{u}_N)$ the collection of all the agents' control sequences but the i-th one.

At every time step k each agent aims at minimizing its own cost function:

$$J_i(\mathbf{u}, x_k) = x_{\tau|k}^\top Q_{i,T} x_{\tau|k} + \sum_{\tau=0}^{T-1} l_i(x_{\tau|k}, u_{\tau|k,i}) \tag{3}$$

with respect to $\mathbf{u}_i$ given $\mathbf{u}_{-i}$, where $x_{\tau|k}$, $\tau = 0, \dots, T$ denotes the predicted state at time $k + \tau$ according to the dynamics in (1) and the initial condition $x_{0|k} = x_k$, with stage cost $l_i(x, u) = x^\top Q_i x + u^\top R_i u$ defined by the agent's weight matrices $Q_i \in \mathbb{S}^{n_x}_{\succcurlyeq 0}$, and $R_i \in \mathbb{S}^{n_{u_i}}_{\succ 0}$, and terminal cost defined by $Q_{i,T} \in \mathbb{S}^{n_x}_{\succcurlyeq 0}$. Note that J_i depends on the collective control sequences $\mathbf{u}$, since all $\mathbf{u}_i$ concur, in general, to evolve the state trajectory $\{x_{\tau|k}\}_{\tau=0}^{T}$ over the prediction interval $[k, \dots, k+T]$. Note that we consider here quadratic stage costs for computational tractability, although our arguments remain compatible with any stage cost so that the mapping from $\mathbf{u}_i$ to $J_i(\mathbf{u}, x_k)$ is convex for all $\mathbf{u}_{-i}$, x_k.

We summarize the collection of the state constraints $x_{\tau|k} \in \mathcal{X}$, $\tau = 1, \dots, T$, in the following shared linear inequalities

$$\mathcal{A}_x \mathbf{u} \leq \boldsymbol{b}_x + \mathcal{C} x_k \tag{4}$$

and the input constraints $u_{\tau|k,i} \in \mathcal{U}_i$, $\tau = 0, \dots, T-1$, $i = 1, \dots, N$ in the following local linear inequalities

$$\mathcal{A}_{u,i} \mathbf{u}_i \leq \boldsymbol{b}_{u,i}, \quad i = 1, \dots, N \tag{5}$$

Moreover, as we will detail in Section III, we will further introduce a shared constraint on the first control input $u_{0|k}$ of the form

$$S\mathbf{u} \in \mathcal{P}(x_k) \tag{6}$$

where $\mathcal{P}(x_k) \subset \mathbb{R}^{n_u}$ is a polytope to be designed as a function of x_k and $S = [I\ 0\ \dots\ 0] \in \mathbb{R}^{n_u \times Tn_u}$ selects the

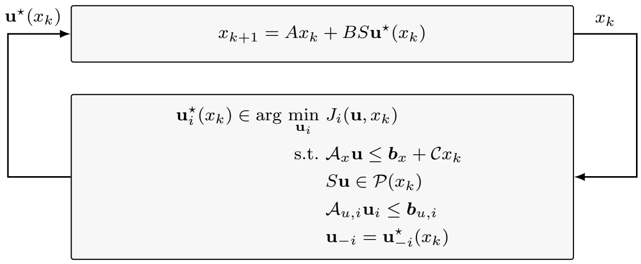

Fig. 1: The closed-loop scheme, consisting of the LTI system (1) controlled in feedback by the GNE of the RHG in (7).

first control move $u_{0|k}$ from the collective control sequence $\mathbf{u}$. The design of $\mathcal{P}(x_k)$ will be crucial for establishing our main closed-loop stability result.

A. Receding-horizon game

By assuming that each agent has a local copy of the dynamics in (1) and the measurement of the full state x_k at each time k, we define the following linear-quadratic (LQ) generalized Nash equilibrium problem (GNEP) for the agents' control sequences $\mathbf{u}_i$, $i = 1, \ldots, N$:

$$\mathbf{u}_i^\star(x_k) \in \arg\min_{\mathbf{u}_i} \; J_i(\mathbf{u}, x_k) \tag{7a}$$
$$\text{s.t. } \mathcal{A}_x \mathbf{u} \le \boldsymbol{b}_x + \mathcal{C} x_k \tag{7b}$$
$$S\mathbf{u} \in \mathcal{P}(x_k) \tag{7c}$$
$$\mathcal{A}_{u,i}\mathbf{u}_i \le \boldsymbol{b}_{u,i} \tag{7d}$$
$$\mathbf{u}_{-i} = \mathbf{u}_{-i}^\star(x_k). \tag{7e}$$

In the considered receding horizon framework, at every time step $k \in \mathbb{N}$ the agents compute a solution to the GNEP (7), yielding a collection of optimal input sequences $\mathbf{u}^\star(x_k)$, and then apply the first element of each $\mathbf{u}_i^\star(x_k)$, i.e., $u_{0|k,i}^\star(x_k)$, to the LTI system in (1). Note that, since $\mathbf{u}_k^\star(x_k)$ depends on the current state x_k, in spite of the *open-loop* information structure of the dynamic game, the overall control strategy is a *feedback* policy, i.e.,

$$u_k = S\mathbf{u}^\star(x_k). \tag{8}$$

This results in the closed-loop system illustrated in Fig. 1:

$$x_{k+1} = Ax_k + BS\mathbf{u}^\star(x_k). \tag{9}$$

Remark 1: The way (distributed, semi-decentralized, or centralized) the GNE is computed at every iteration $k \in \mathbb{N}$ is not in the focus of this paper. We tacitly assume there is a procedure in place that allows the agents to compute an equilibrium at every step k. In this paper, we will use the centralized mixed-integer linear programming (MILP) formulation to solve LQ-GNEPs described in [9].

Without any further assumption on the formulation of the RHG in (7), there is no guarantee that the feedback loop (9) is asymptotically stable. Moreover, we need to guarantee that a GNE in (7) exists at every time step $k \in \mathbb{N}$. The purpose of this work is to design the shared constraint (7c) so that the closed-loop dynamics in (9) converges exponentially to the origin and at least a GNE $\mathbf{u}^\star(x_k)$ is guaranteed to exist at every $k \in \mathbb{N}$; in particular, we want to avoid that an inaccurate choice of $\mathcal{P}$ even leads to the local and shared constraints defining an empty set. The choice of $\mathcal{P}$ will be based on control-invariant sets, as detailed in the next section.

III. Polyhedral control Lyapunov functions

As a key consequence of Assumption 1, we note that one can uniquely associate with the pair (A, B) and the constraints in (2) the so-called MCPI set [22], which is λ-contractive according to the following definition:

Definition 1: Given $\lambda \in (0, 1]$, a set $\mathcal{X}^\lambda \subseteq \mathcal{X}$ is λ-contractive control positive invariant (λ-CPI) for the LTI system (1) if for every $x \in \mathcal{X}^\lambda$ there exists an input $u \in \mathcal{U}$ such that $\frac{1}{\lambda}(Ax + Bu) \in \mathcal{X}^\lambda$.

In the case $\lambda = 1$, $\mathcal{X}^1$ is also called control positive invariant (CPI). Thus, for a given λ, the maximal λ-contractive (λ-MCPI) set $\mathcal{X}_\infty^\lambda \subseteq \mathcal{X}$ turns out to coincide with the union of all λ-CPI sets for (1). Note that the union of λ-CPI sets is itself a λ-CPI set.

The way we propose to design the set $\mathcal{P}$ in (7c) relies on $\mathcal{X}_\infty^\lambda$ and the associated contractive property. Computing the λ-MCPI set usually relies on viability-based iterative procedures—see, e.g., [23], [20], [24]. However, the adopted scheme is not guaranteed to converge in a finite number of iterations [22]. To this end, we properly select λ and approximate $\mathcal{X}_\infty^\lambda \approx \mathcal{X}_{k_F}^\lambda$ for some (large enough) k_F, where $\mathcal{X}_{k_F}^\lambda$ is a polytope obtained by the intersection of n_F half-spaces, with the origin in its interior, cf. [17]. With a slight abuse of notation, from now on we identify the λ-MCPI for (1) w.r.t. (2) with:

$$\mathcal{X}_\infty^\lambda = \{x \in \mathbb{R}^{n_x} \mid Fx \le \mathbf{1}_{n_F}\} \tag{10}$$

where $F \in \mathbb{R}^{n_F \times n_x}$ and $\max_{i=1,\ldots,n_F}(F_i^\top x) > 0$ for all $x \in \mathbb{R}^{n_x} \backslash \{0\}$ since a polytope is a bounded set. A polyhedral inner approximation as the one in (10) can be obtained, indeed, in a finite number of iterations and with arbitrary precision [20, Proposition 5.9]. Note that λ is chosen to be strictly greater than the spectral radius of the uncontrollable part of the pair (A, B), if such part exists.

In view of its structure in (10), we recall that the set $\mathcal{X}_\infty^\lambda$ represents the DoA of the following second-order[1] PCLF:

$$V_\lambda(x) = \left(\max_{i=1,\ldots,n_F}(F_i^\top x)\right)^2. \tag{11}$$

Although non-smooth, $V_\lambda(\cdot)$ has been exploited for the stabilization of LTI systems with affine state and input constraints in various ways. For instance, it can be used (modulo proper weighting) as terminal cost in the MPC framework for the stabilization of the origin. In the same context, $\mathcal{X}_\infty^\lambda$ denotes the DoA of the controller, as well as the terminal set. Unlike available approaches, in this work we rely on the following two properties of the PCLF $V_\lambda(\cdot)$ in (11):

[1] The results of this paper can be immediately replicated for the polyhedral Lyapunov function $V_\lambda(x) = \max_{i=1,\ldots,n_F}(F_i^\top x)$ as well.

Lemma 1 ([17]): Let the λ-MCPI set $\mathcal{X}_\infty^\lambda$ be as in (10) and the associated PCLF $V_\lambda(\cdot)$ as in (11). Then:

(i) $\forall x \in \mathcal{X}_\infty^\lambda,\ \exists u \in \mathcal{U}:\ 0 \le V_\lambda(Ax+Bu) \le \lambda^2 V_\lambda(x)$,
(ii) $\exists c_1,\, c_2 > 0 : c_1\|x\|^2 \le V_\lambda(x) \le c_2\|x\|^2$.

While Lemma 1(i) stems from the contractivity of $\mathcal{X}_\infty^\lambda$ and the homogeneity of the corresponding Minkowski function $\Psi(x) = \max_{i=1,\dots,n_F}(F_i^\top x)$, the second part follows from the upper and lower bounds of $\Psi(\cdot)$ itself. Such bounds are based on the state norm $\|x\|$, which can in turn be derived using $\|F\|$ and the standard definition of the Minkowski function.

IV. Stability analysis of the RHG

In view of the discussion above, for a given $\lambda \in (0,1)$ we then propose to rewrite the set inclusion in (7c) starting from the condition stated in Lemma 1(i). For any $x_k \in \mathcal{X}_\infty^\lambda$ there indeed exists $u_{0|k} = S\mathbf{u}$ such that:

$$\left(\max_{i=1,\dots,n_F}(F_i^\top(Ax_k + BS\mathbf{u}))\right)^2 \le \lambda^2\left(\max_{i=1,\dots,n_F}(F_i^\top x_k)\right)^2.$$

Due to the fact that $\max_{i=1,\dots,n_F}(F_i^\top x) > 0$, for all $x \in \mathbb{R}^{n_x}\setminus\{0\}$, the inequality above is therefore equivalent to:

$$\max_{i=1,\dots,n_F}(F_i^\top(Ax_k + BS\mathbf{u})) \le \lambda \max_{i=1,\dots,n_F}(F_i^\top x_k)$$

and thus to:

$$FBS\mathbf{u} \le \lambda\mathbf{1}_{n_F}\max_{i=1,\dots,n_F}(F_i^\top x_k) - FAx_k. \tag{12}$$

Before stating and proving our main stability result involving the closed-loop dynamics in (9), let us prove the following preliminary result.

Lemma 2: Let $x_k \in \mathcal{X}_\infty^\lambda$ with (7c) formulated as in (12). Then, the GNEP in (7) admits at least one GNE.

Proof: Consider the collective feasible set

$$\Omega(x_k) \triangleq \left\{\mathbf{u} :\ \begin{aligned} &\mathcal{A}_{u,i}\mathbf{u}_i \le \boldsymbol{b}_{u,i},\ i=1,\dots,N \\ &\mathcal{A}_x\mathbf{u} \le \boldsymbol{b}_x + \mathcal{C}x_k \\ &FBS\mathbf{u} \le \lambda\mathbf{1}_{n_F}\max_{i=1,\dots,n_F}(F_i^\top x_k) - FAx_k \end{aligned}\right\} \tag{13}$$

and the pseudo-gradient mapping $\mathcal{F} : \mathbb{R}^{Tn_u} \to \mathbb{R}^{Tn_u}$:

$$\mathcal{F}(\mathbf{u}; x_k) = \mathrm{col}(\nabla_{\mathbf{u}_i} J_i(\mathbf{u}, x_k))_{i=1}^N. \tag{14}$$

Clearly, $\Omega(x_k)$ is convex and compact. To prove that it is also nonempty, consider the following input sequence $\hat{\mathbf{u}}$ generated as follows:

$$\hat{\mathbf{u}} = \left[\hat{u}_{0|k}^\top, \dots, \hat{u}_{T-1|k}^\top\right]^\top.$$

Since $x_k \in \mathcal{X}_\infty^\lambda$, by Lemma 1 there exists a $\hat{u}_{0|k} \in \mathcal{U}$ such that $x_{1|k} = Ax_k + B\hat{u}_{0|k} \in \lambda\mathcal{X}_\infty^\lambda \subset \mathcal{X}_\infty^\lambda \subseteq \mathcal{X}$. By iterating this argument, we can find the remaining values $\hat{u}_{1|k}, \dots, \hat{u}_{T-1|k}$. Thus, the input sequence $\hat{\mathbf{u}}$ satisfies all the constraints in (13), which proves the non-emptiness of $\Omega(x_k)$.

Moreover, $\mathcal{F}$ is affine, and hence continuous, in $\mathbf{u}$. Consequently, by the Hartman–Stampacchia theorem, the following variational inequality problem, $\mathrm{VI}(\mathcal{F}(\cdot; x_k), \Omega(x_k))$, i.e.,

$$\text{Find } \mathbf{u}^\star \in \Omega(x_k) \text{ s.t. } \mathcal{F}(\mathbf{u}^\star; x_k)^\top(\mathbf{u} - \mathbf{u}^\star) \ge 0,\ \forall \mathbf{u} \in \Omega(x_k) \tag{15}$$

admits at least one solution $\mathbf{u}^\star$. As every function J_i in the GNEP (7) is convex w.r.t. $\mathbf{u}_i$ for all $\mathbf{u}_{-i}$ and x_k, by [21, Theorem 5], every solution to (15) is a GNE of (7). ■

Theorem 1: Let $x_0 \in \mathcal{X}_\infty^\lambda$. Then, at every step $k \in \mathbb{N}$ a GNE $\mathbf{u}^\star(x_k)$ of the RHG in (7) with (12) in place of (7c) exists and the control strategy (8) exponentially stabilizes the LTI system in (1) to the origin.

Proof: Since $x_0 \in \mathcal{X}_\infty^\lambda$, by Lemma 2 we have that the generalized Nash game in (7) with (12) in place of (7c) admits a GNE.

Now, assume that the RHG admits at least a GNE at the generic time step k. Due to the constraints (12) and the feedback policy (8) one readily obtains:

$$\max_{i=1,\dots,n_F}(F_i^\top x_{k+1}) \le \lambda \max_{i=1,\dots,n_F}(F_i^\top x_k) < 1 \tag{16}$$

thereby implying that $x_{k+1} \in \mathcal{X}_\infty^\lambda$. Thus, at time step $k+1$ the generalized Nash game at hand meets the conditions described in Lemma 2 and, hence a GNE exists again. By induction on k, we can therefore claim the recursive well-posedness on the control law (8) for all $k \in \mathbb{N}$.

To prove exponential stabilization to the origin, we first want to show by induction that the following relation holds true:

$$V_\lambda(x_k) \le \lambda^{2k} V_\lambda(x_0),\ \forall k > 0. \tag{17}$$

Indeed, this relation holds for $k = 1$, since the coupling constraints (12) correspond to the condition in Lemma 1(i). Now, assume that (17) holds true for some $k > 0$. By the recursive feasibility proved above, a GNE $\mathbf{u}^\star(x_k)$ exists and is constrained to satisfy (12), which forces $F_i^\top x_{k+1} \le \lambda \max_j F_j^\top x_k,\ \forall i$. Hence, recalling the definition of $V_\lambda(x_{k+1})$ in (11), we obtain:

$$V_\lambda(x_{k+1}) \le \lambda^2 V_\lambda(x_k) \le \lambda^{2(k+1)} V_\lambda(x_0) \tag{18}$$

which proves the desired induction. The proof then concludes by the lower bound on $V_\lambda(x_k)$ and upper bound on $V_\lambda(x_0)$ from Lemma 1(ii), together with (17):

$$c_1\|x_k\|^2 \le V_\lambda(x_k) \le \lambda^{2k} V_\lambda(x_0) \le \lambda^{2k} c_2\|x_0\|^2$$

which yields

$$\|x_k\| \le \lambda^k \sqrt{\frac{c_2}{c_1}}\,\|x_0\| \tag{19}$$

and, hence, proves exponential convergence to the origin. ■

Note that the stability result in Theorem 1 is independent of the choice of the cost functions J_i in (3) (and, in particular, of the terminal costs), as long as the resulting mapping from $\mathbf{u}_i$ to $J_i(\mathbf{u}, x_k)$ is convex for all $\mathbf{u}_{-i}$ and x_k. More specifically, if the cost functions J_i are also defined so that the resulting pseudo-gradient $\mathcal{F}$ in (14) is strongly monotone, a variational GNE exists and is unique at each time step. This is an easy condition to check in our quadratic setting (3).

TABLE I: Model and receding-horizon game parameters

Parameter	Value [a.u.]
Mass M_1	3
Mass M_2	1
Spring constants k_1, k_2	0.18
Spring constant k	0.5
Friction coeff. β_1	1.5
Friction coeff. β_2	1.0
Sampling Time T_s	0.2
Lower bound x_{min}	$-[15\ 10\ 15\ 10]^\top$
Upper bound x_{max}	$-x_{\text{min}}$
Lower bound u_{min}	$-[100\ 100]^\top$
Upper bound u_{max}	$-u_{\text{min}}$
Horizon length T	2
Stage cost coeff. Q_1	diag[1, 1, 0, 0]
Stage cost coeff. Q_2	diag[0, 0, 1, 1]
Stage cost coeff. $R_1 = R_2$	1
Decay rate λ	0.8, 0.9, 0.9999
Approximation parameter ϵ	0.01

V. Numerical experiments

To corroborate our theoretical findings we run numerical simulations on the mechanical system illustrated in Fig. 2, with numerical parameters reported in Tab. I. The latter also reports all the design choices we made for the experiments. The approximation parameter ϵ determines both the decay rate used in the viability-based, iterative process and its termination condition.

The underlying system consists of two masses M_1 and M_2 connected to walls by two linear springs with identical constants $k_1 = k_2$. The masses are additionally connected to each other by another linear spring with parameter k and, due to the interaction with the surface on which they slide horizontally, are subject to viscous friction with coefficients β_1 and β_2. Each mass is controlled by an agent i, applying its force input f_i to set the position y_i of the center of gravity of M_i to its origin. All springs are initialized to their natural lengths.

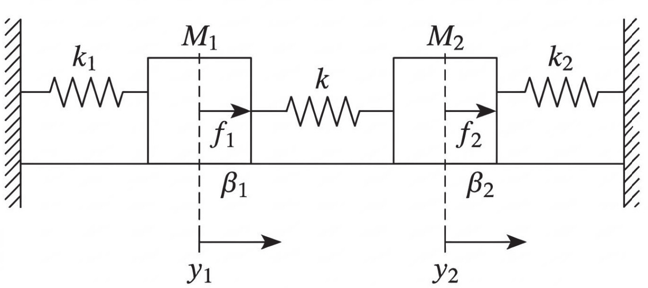

Fig. 2: Schematic representation of the two-mass system considered.

For consistency of notation with Section II, we define the state vector $x = [y_1\ \dot{y}_1\ y_2\ \dot{y}_2]^\top$, the collective input $u = [f_1\ f_2]^\top$ and the following LTI system in continuous time:

$$\dot{x} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ -\frac{(k+k_1)}{M_1} & -\frac{\beta 1}{M_1} & \frac{k}{M_1} & 0 \\ 0 & 0 & 0 & 1 \\ \frac{k}{M_2} & 0 & -\frac{(k_2+k)}{M_2} & -\frac{\beta_2}{M_2} \end{bmatrix} x + \begin{bmatrix} 0 & 0 \\ \frac{1}{M_1} & 0 \\ 0 & 0 \\ 0 & \frac{1}{M_2} \end{bmatrix} u. \tag{20}$$

We obtain a discrete-time prediction model by discretizing (20) with sampling time T_s and zero-order-hold (ZOH) on the inputs. The resulting state variables and control input of each agent are subject to the following box constraints:

$$x_{\text{min}} \leq x_k \leq x_{\text{max}} \tag{21a}$$

$$u_{\text{min}} \leq u_k \leq u_{\text{max}} \tag{21b}$$

where $x_{\text{min}},\ x_{\text{max}} \in \mathbb{R}^4$ and $u_{\text{min}},\ u_{\text{max}} \in \mathbb{R}^2$.

We approximate the λ-MCPI set for the aforementioned discretized dynamics subject to (21) using the backward viability iteration described in [20, Proposition 5.9], with contraction factor $\frac{\lambda}{1+\epsilon}$. The iteration is implemented using the Multi-Parametric Toolbox for MATLAB [25].

To evaluate the behavior of the closed-loop system given in (9), we simulate it for a sequence of 100 initial conditions x_0 on the boundary of the computed λ-contractive set, for $\lambda = 0.9$. We consider a simulation horizon of 90 time steps, at each of which, the GNE of the GNEP (7), with the constraint (7c) replaced by inequality (12), is computed using the MILP formulation described in [9].

At every single simulated time step k we compute the ratio $\frac{\|x_k\|}{\|x_0\|}$. We then find the maximum, the minimum and the mean value of this ratio at each time step k among all sampled trajectories. Moreover, in view of inequality (19) we calculate the reference quantity $r_k = \lambda^k \sqrt{\frac{c_2}{c_1}}$ and present the results in Fig. 3. In the same figure, we also report the empirical relative frequency of the ratio $\frac{\|x_k\|}{\|x_0\|}$ across all sampled trajectories at each time step k. More precisely, for every k, the range between the observed minimum and maximum values is partitioned into bins, and the relative frequency associated with each bin is defined as the number of trajectories whose normalized state norm falls into that bin at time step k, divided by the largest bin count at the same time step. As expected, inequality (19) is satisfied for every initial condition that renders the GNEP well-posed and at every time step k, without a single exception.

Furthermore, we examine the effect of the decay rate λ on the size of the approximated λ-MCPI set and the closed-loop performance of the RHG. For each of the three values of λ reported in Table I we compute the set $\mathcal{X}_\infty^\lambda$ and we define its relative volume as $\nu_\lambda = \text{vol}(\mathcal{X}_\infty^\lambda)/\,\text{vol}(\mathcal{X})$. Subsequently, starting from a common nonzero initial condition lying strictly inside the intersection of the resulting sets, we simulate the closed-loop system given in (9) for each value of λ. We approximate each player's infinite-horizon closed-loop cost J_i^∞ by summing the stage costs along the simulated trajectory. We denote the resulting approximations by $\widehat{J}_i^\infty$ and we summarize the results in Table II.

We observe that increasing λ enlarges the computed λ-MCPI set. To some extent, smaller values of λ yield both

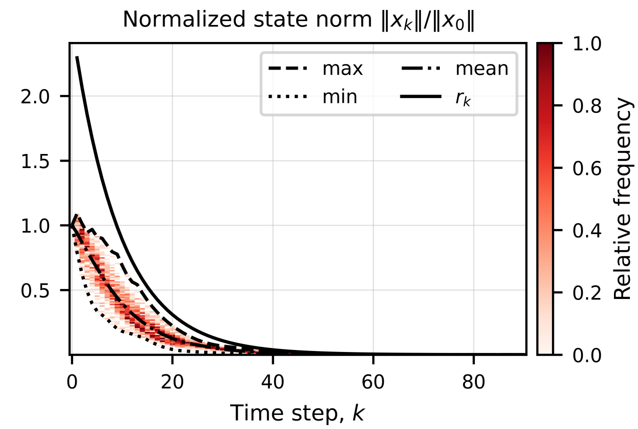

Fig. 3: The maximum, the minimum and the mean value and the relative frequency of the ratio $\frac{\|x_k\|}{\|x_0\|}$, at fixed time step k, along the sampled trajectories.

TABLE II: Relative volumes of the computed λ-MCPI sets and approximate infinite-horizon closed-loop costs for different decay rates.

λ	$\nu_\lambda[\%]$	$\widehat{J}_1^\infty[a.u.]$	$\widehat{J}_2^\infty[a.u.]$	$(\widehat{J}_1^\infty + \widehat{J}_2^\infty)[a.u.]$
0.8	27.26	10189.71	1369.35	11559.05
0.9	89.95	2945.23	568.86	3514.08
0.9999	99.56	935.50	824.79	1760.28

smaller relative volumes ν_λ and larger aggregate infinite-horizon cost approximations $\widehat{J}_1^\infty + \widehat{J}_2^\infty$. However, the individual cost approximations $\widehat{J}_1^\infty$ and $\widehat{J}_2^\infty$ do not follow the same ordering.

VI. Conclusion

For LTI discrete-time multi-agent systems subject to polytopic state and input constraints, we have proposed a game-theoretic MPC scheme to stabilize the system at the origin. Regardless of the choice of the performance indices optimized by the agents and the particular equilibrium selected in case of multiple GNEs, exponential convergence is guaranteed by imposing affine shared constraints on all agents, that are derived from the use of maximal $\lambda-$contractive MCPI sets. Current research is devoted to extending the approach to handle uncertain dynamics.

References

[1] E. Benenati and S. Grammatico, "Linear-quadratic dynamic games as receding-horizon variational inequalities," *IEEE Transactions on Automatic Control*, vol. 70, no. 1, pp. 1–16, 2025.

[2] S. Le Cleac'h, M. Schwager, and Z. Manchester, "Algames: A fast augmented Lagrangian solver for constrained dynamic games," *Autonomous Robots*, vol. 46, no. 1, pp. 201–215, 2022.

[3] R. Spica, E. Cristofalo, Z. Wang, E. Montijano, and M. Schwager, "A real-time game theoretic planner for autonomous two-player drone racing," *IEEE Transactions on Robotics*, vol. 36, no. 5, pp. 1389–1403, 2020.

[4] S. Hall, L. Guerrini, F. Dörfler, and D. Liao-McPherson, "Game-theoretic model predictive control for modelling competitive supply chains," in *Proc. 8th IFAC Conference on Nonlinear Model Predictive Control*, vol. 58, no. 4, 2024, pp. 8–14.

[5] S. Hall, G. Belgioioso, D. Liao-McPherson, and F. Dörfler, "Receding horizon games with coupling constraints for demand-side management," in *Proc. IEEE Conference on Decision and Control*, 2022, pp. 5288–5294.

[6] S. Hall, D. Liao-McPherson, G. Belgioioso, and F. Dörfler, "Stability certificates for receding horizon games," *IEEE Transactions on Automatic Control*, 2025.

[7] F. Fabiani and A. Simonetto, "Incentives and coevolution: Steering linear dynamical systems with noncooperative agents," *IEEE Transactions on Control of Network Systems*, vol. 11, no. 2, pp. 1067–1079, 2023.

[8] A. Monti, B. Nortmann, T. Mylvaganam, and M. Sassano, "Feedback and open-loop Nash equilibria for LQ infinite horizon discrete-time dynamic games," *SIAM Journal on Control and Optimization*, vol. 62, no. 3, pp. 1417–1436, 2024.

[9] A. Bemporad, "NashOpt: A Python library for computing generalized Nash equilibria," *arXiv preprint arXiv:2512.23636*, 2025. [Online]. Available: https://arxiv.org/abs/2512.23636

[10] S. Hall and A. Bemporad, "Solving multiparametric generalized Nash equilibrium problems and explicit game-theoretic model predictive control," 2025, available on arXiv at https://arxiv.org/abs/2512.05505.

[11] E. Benenati and G. Belgioioso, "The explicit game-theoretic linear quadratic regulator for constrained multi-agent systems," *arXiv preprint arXiv:2512.07749*, 2025.

[12] A. N. Venkat, I. A. Hiskens, J. B. Rawlings, and S. J. Wright, "Distributed MPC strategies with application to power system automatic generation control," *IEEE Transactions on Control Systems Technology*, vol. 16, no. 6, pp. 1192–1206, 2008.

[13] E. Camponogara, D. Jia, B. H. Krogh, and S. Talukdar, "Distributed model predictive control," *IEEE Control Systems Magazine*, vol. 22, no. 1, pp. 44–52, 2002.

[14] A. Alessio and A. Bemporad, "Decentralized model predictive control of constrained linear systems," in *Proc. European Control Conference*, 2007, pp. 2813–2820.

[15] P.-O. Gutman and M. Cwikel, "Admissible sets and feedback control for discrete-time linear dynamical systems with bounded controls and states," *IEEE Transactions on Automatic Control*, vol. 31, no. 4, pp. 373–376, 1986.

[16] S. S. Keerthi and E. G. Gilbert, "Computation of minimum-time feedback control laws for discrete-time systems with state-control constraints," *IEEE Transactions on Automatic Control*, vol. 32, no. 5, pp. 432–435, 1987.

[17] S. Grammatico and G. Pannocchia, "Linear model predictive control based on polyhedral control Lyapunov functions: Theory and applications," arXiv:1208.2429, 2012.

[18] A. Bemporad, F. Borrelli, and M. Morari, "Model predictive control via linear programming—the explicit solution," *IEEE Transactions on Automatic Control*, vol. 47, no. 12, pp. 1974–1985, 2002.

[19] M. Lazar, W. P. M. H. Heemels, S. Weiland, A. Bemporad, and O. Pastravanu, "Infinity norms as Lyapunov functions for model predictive control of constrained PWA systems," in *Proc. 13th Mediterranean Conference on Control and Automation*, 2005, pp. 1006–1011.

[20] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*. Basel, Switzerland: Birkhäuser, 2008.

[21] F. Facchinei and C. Kanzow, "Generalized Nash equilibrium problems," *Annals of Operations Research*, vol. 175, no. 1, pp. 177–211, 2010.

[22] E. C. Kerrigan, "Robust constraint satisfaction: Invariant sets and predictive control," Ph.D. dissertation, University of Cambridge, Cambridge, U.K., Nov. 2000.

[23] F. Blanchini, "Ultimate boundedness control for uncertain discrete-time systems via set-induced Lyapunov functions," *IEEE Transactions on Automatic Control*, vol. 39, no. 2, pp. 428–433, 1994.

[24] S. Miani and C. Savorgnan, "MAXIS-G: A software package for computing polyhedral invariant sets for constrained LPV systems," in *Proc. IEEE Conference on Decision and Control*, Seville, Spain, 2005, pp. 7609–7615.

[25] M. Herceg, M. Kvasnica, C. Jones, and M. Morari, "Multi-Parametric Toolbox 3.0," in *Proc. of the European Control Conference*, Zürich, Switzerland, July 17–19 2013, pp. 502–510.