A Two-stage Mechanism for Prioritized Trajectory Planning in Multi-Agent Systems

Yu-Wen Chen, Can Kizilkale, Murat Arcak

Abstract— **In multi-agent systems with coupled objectives and/or constraints, agents may misreport information to achieve individual gains. This issue is exacerbated when agents possess local decision-making power, such as in multi-agent trajectory planning, where the increased autonomy amplifies individual benefits at the expense of a higher social cost. To overcome this problem, we leverage the Vickrey-Clarke-Grove (VCG) framework and propose a strategyproof, two-stage mechanism. We further extend this mechanism to prioritized planning and prevent agents from manipulating their priority.**

I. INTRODUCTION

Trajectory planning has been studied extensively, leading to various methodologies, such as the A* algorithm [1], Rapidly-exploring Random Tree and Probabilistic Road Map methods [2], optimal control approaches [3], [4], evolutionary algorithms [5], diffusion models [6], geometric methods [7], and hierarchical decompositions [8]–[10]. An important research theme now is to extend these methods to Multi-Agent Systems (MASs), [11]. In MASs additional factors must be considered, such as agent interactions, fairness in handling coupling constraints, and mismatches between agent- and system-level objectives. In a typical MAS trajectory planning, each agent has objectives (often private) and constraints (possibly coupled with the other agents), while the system designer has a global objective, such as minimizing the sum of agent objectives.

Integrating individual objectives into a system-level goal falls within the scope of multi-objective optimization [12]. Some approaches employ evolutionary algorithms to explore the entire Pareto-optimal front [13], while others focus on specific solutions within this front by combining objectives using prioritization or weightings. This widely adopted method is known as prioritized planning, where either a weighted sum of agents' objective functions is optimized or a priority order is assigned to optimize each function sequentially. The significance of determining an appropriate priority order is highlighted in [14].

The solution techniques for MAS trajectory planning can be classified as centralized [14]–[17] or distributed [18]–[20]. In centralized methods, a central planner gathers information from all agents to solve a system-level optimization problem and communicates the planned trajectory back to the agents. In distributed methods, agents share local information with their neighbors iteratively, following predefined protocols, until convergence is achieved. For rational agents who selfishly select strategies that maximize their own utility, the problem can be modeled as optimizing the outcome at equilibrium [14], [15].

In this paper, we consider MAS trajectory planning where there is a central planner, but agents have local decision-making power. Instead of strictly adhering to a trajectory received from the central planner, the agent views it as a recommendation and fine-tunes it to minimize its own cost and to satisfy its local constraints.

An important consideration in this setup is to ensure that agents have no incentive to misrepresent their private information. One way to achieve this is to introduce appropriate pricing mechanisms. For instance, in second-price auctions, each agent submits their valuation of an item and the highest bidder wins, but it pays the second-highest bid. Thus, bidders cannot benefit from misrepresenting their true valuation. This is an example of the Vickrey–Clarke–Groves (VCG) mechanism [21], which ensures that reporting truthfully is a weakly dominant strategy.

Our contributions are two-fold. First, we propose a two-stage mechanism that adapts the VCG framework to trajectory planning when agents have local decision-making power. In the first stage, based on the information reported by agents, a reference signal is generated and sent to all agents. In the second stage, agents locally determine their optimal decisions and pay according to prescribed payment rules. We show that, with properly designed reference signals and payment functions, this mechanism is strategyproof; that is, it ensures that truthful reporting of objective functions and constraints is a weakly dominant strategy. Next, we extend the two-stage mechanism to prioritized planning.

References [16], [17] also make use of the VCG framework, but focus on the scenario where the planner's decisions are binding and must be followed by all players. Another related reference is [14], which shows the advantage of determining the priority order based on agents' objective functions. However, when priority order is based on the reported information, agents may misreport to influence the planner's decision – either directly by distorting the optimization problem, or indirectly by securing a higher priority which further impacts the decision-making process. While VCG payments can resolve the former issue, we propose additional designs to deal with the latter.

The remainder of the paper is structured as follows. Section II provides an overview of VCG mechanisms. Section III presents the two-stage mechanism and proves the dominance of the truthful reporting strategy. Section IV extends the

Y.-W. Chen, C. Kizilkale and M. Arcak are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA. {yuwen_chen, cankizilkale, arcak}@berkeley.edu. This work was supported by the National Science Foundation grant CNS-2135791.

mechanism to prioritized planning. Section V illustrates the results on a case study of constrained Linear Quadratic Regulator. Finally, Section VI gives conclusions.

II. An Overview of VCG Mechanisms

In game theory, a mechanism is a set of rules that governs how participants make decisions and receive outcomes to achieve efficiency and fairness. In scenarios such as goods allocation through auctions or decisions regarding public projects, participants may strategically misrepresent their private preferences. VCG mechanisms provide a strategyproof way to implement efficient allocations for such problems in quasi-linear environments (i.e., environments where participants' preferences are linear with respect to money), ensuring that participants truthfully reveal their preferences.

Consider a set of players $i = 1, \cdots, N$, and a set of possible outcomes, denoted by X. Each player i has a valuation function $v_i : X \to \mathbb{R}_+$ and reports it as $\tilde{v}_i$, which is not necessarily equal to v_i since the player may misreport. For every outcome, the player is also assigned a payment p_i. Based on the reported functions, the outcome is given by

$$x^* = \arg\max_x \sum_{i=1}^{N} \tilde{v}_i(x), \tag{1}$$

and the payment is

$$p_i = \sum_{j \neq i} \tilde{v}_j(x^*) + h_i(\tilde{v}_{-i}), \tag{2}$$

where $h_i(\tilde{v}_{-i})$ is a function that depends on, $\tilde{v}_{-i}$, the reported functions of all players except i. Thus, player i seeks to maximize its utility, $v_i(x) + p_i$. The optimization problem for agent i can then be expressed as

$$\max_{\tilde{v}_i} \; v_i\left(x^*(\tilde{v}_i, \tilde{v}_{-i})\right) + \sum_{j \neq i} \tilde{v}_j\left(x^*(\tilde{v}_i, \tilde{v}_{-i})\right) + h_i(\tilde{v}_{-i}). \tag{3}$$

Since $h_i(\tilde{v}_{-i})$ is independent of $\tilde{v}_i$, it can be excluded from the optimization. Therefore, reporting $\tilde{v}_i = v_i$ aligns (1) with (3), ensuring that x^* maximizes player i's utility (3), making truthful reporting the best response for all players. A common choice of $h_i(\tilde{v}_{-i})$ is $-\max_x \sum_{j \neq i} \tilde{v}_j(x)$, in which case p_i is player i's contribution to the social cost.

III. A Strategyproof Two-stage Mechanism for MAS Trajectory Planning

A. Problem Setting

Consider a set of N agents, and let $[N] = \{1, \cdots, N\}$ represent the index set of agents. For each $i \in [N]$, agent i satisfies the discrete-time dynamics:

$$x_i[k+1] = f_i\left(x_i[k], u_i[k]\right), \tag{4}$$

where $x_i[k] \in \mathbb{R}^{n_i}$ and $u_i[k] \in \mathbb{R}^{m_i}$ denote the state and input for agent i at time instant $k \in \mathbb{N}_0 = \{0, 1, 2, \dots\}$, respectively. We denote $x_i : \mathbb{N}_0 \to \mathbb{R}^{n_i}$ and $u_i : \mathbb{N}_0 \to \mathbb{R}^{m_i}$ as the state and input trajectory for agent i, respectively. Moreover, we denote the composite state trajectories of all agents as $x = (x_1, \cdots, x_N)$ and the composite input trajectories of all agents as $u = (u_1, \cdots, u_N)$.

Each agent i has a nonnegative cost function c_i evaluating each pair (x, u) by $c_i(x, u)$. The central planner poses a price, $p_i(x, u)$, for each agent i based on the state and input trajectories. Each agent i has coupling constraints, $\bar{g}_i(x, u) \preceq 0$. Therefore, each agent i solves the optimization problem:

$$\begin{aligned} \arg\min_{x_i, u_i} \quad & c_i(x, u) + p_i(x, u) \\ s.t. \quad & \text{the dynamics (4) is satisfied} \\ & \bar{g}_i(x, u) \preceq 0. \end{aligned} \tag{5}$$

Let $y_i = (x_i, u_i)$ denote all decision variables of agent i, and let $y = (y_1, \cdots, y_N)$ represent the composite decision variables of all agents. We use the subscript $-i$ to indicate the variables of all agents except for agent i, *e.g.*, $y_{-i} = (y_1, \cdots, y_{i-1}, y_{i+1}, \cdots, y_N)$. Thus, (5) can be rewritten as

$$\begin{aligned} \arg\min_{y_i} \quad & c_i(y_i, y_{-i}) + p_i(y_i, y_{-i}) \\ s.t. \quad & g_i(y_i, y_{-i}) \preceq 0. \end{aligned} \tag{6}$$

B. Proposed Two-stage Mechanism

In the following, we propose a strategyproof two-stage mechanism designed for MASs where agents possess local decision-making power, described in Mechanism 1. Compared with the VCG mechanism discussed in Section II, Mechanism 1 introduces the additional step (ii) where agents locally solve for the optimal trajectories. This accounts for the local decision-making power. From an optimization perspective, in the VCG mechanism, agent i has the decision variable $\tilde{c}_i(\cdot)$, whereas in Mechanism 1, it has decision variables $\tilde{c}_i(\cdot)$ and y_i, involving a larger decision space. In the following, we show that reporting truthfully is a weakly dominant strategy for Mechanism 1.

Theorem 1. *Mechanism 1 is strategyproof.*

Proof. The minimum cost player i can achieve by changing its decision y_i is

$$\min_{(y_i, \tilde{y}^*_{-i}(\tilde{c}, \tilde{g})) \in \mathcal{F}_i} c_i\left(y_i, \tilde{y}^*_{-i}(\tilde{c}, \tilde{g})\right) + p_i\left(y_i, \tilde{y}^*_{-i}(\tilde{c}, \tilde{g})\right) \tag{9}$$

$$\geq \min_{y \in \mathcal{F}_i} c_i(y) + \sum_{j \neq i} \tilde{c}_j(y) - h_i(\tilde{c}_{-i}, \tilde{g}_{-i}), \tag{10}$$

$$= c_i\left(\tilde{y}^*(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})\right) + p_i\left(\tilde{y}^*(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})\right), \tag{11}$$

where $\mathcal{F}_i = \{y : g_i(y) \preceq 0, \tilde{g}_j(y) \preceq 0, \forall j \neq i\}$ and $\tilde{y}^*(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})$ is the reference signal when agent i reports truthfully, and (11) is the total cost of agent i. Thus, reporting $\tilde{c}_i = c_i$ and $\tilde{g}_i = g_i$ is a weakly dominant strategy. ∎

Note that $h_i(\tilde{c}_{-i}, \tilde{g}_{-i})$ can be selected as

$$\begin{aligned} h_i(\tilde{c}_{-i}, \tilde{g}_{-i}) = \min_y \; & \sum_{j \neq i} \tilde{c}_j(y) \\ s.t. \; & \tilde{g}_{-i}(y) \preceq 0, \end{aligned} \tag{12}$$

which mimics the design of externality cost described in (3). The term $I_{\tilde{g}_j}\left(y_i, \tilde{y}^*_{-i}(\tilde{c}, \tilde{g})\right)$ describes that if agent i, after

Mechanism 1 with local decision-making power

(i) Each agent $i \in [N]$ reports its cost function $\tilde{c}_i$ and its constraint function $\tilde{g}_i$, which are not necessarily truthful. Then, based on the reported functions $\tilde{g} = (\tilde{g}_1, \cdots, \tilde{g}_N)$ and $\tilde{c} = (\tilde{c}_1, \cdots, \tilde{c}_N)$ for all i, the planner determines a reference signal

$$\tilde{y}^*(\tilde{c}, \tilde{g}) = \arg\min_{y} \sum_{i=1}^{N} \tilde{c}_i(y) \quad s.t. \quad \tilde{g}(y) \preceq 0, \tag{7}$$

and broadcast it to the agents.

(ii) Upon receiving the reference signal from the planner, each agent i locally decides its optimal trajectory y_i^* by (6) and pays $p_i\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right)$ which is given by,

$$p_i(y_i, y_{-i}) = \sum_{j \neq i} \tilde{c}_j(y_i, y_{-i}) + I_{\tilde{g}_j}(y_i, y_{-i}) - h_i(\tilde{c}_{-i}, \tilde{g}_{-i}), \tag{8}$$

where $h_i(\tilde{c}_{-i}, \tilde{g}_{-i})$ is an arbitrary function based on information reported by others, and

$$I_f(y_i, y_{-i}) = \begin{cases} 0 & \text{if } f(y_i, y_{-i}) \preceq 0 \\ \infty & \text{otherwise.} \end{cases}$$

receiving the reference signal, decides not to follow it and causes a violation, then agent i is responsible for paying an infinity cost.

IV. A Strategyproof Two-stage Mechanism for Prioritized Planning

In (7), the planner adopts a simple approach by summing over agents' objectives as a system-level goal. However, the relative importance of agents' objectives is often critical in MAS applications, which gives rise to prioritized planning. As mentioned in Section I, prioritized planning typically follows two approaches: (i) optimizing a weighted sum of agents' objective functions or (ii) assigning a priority order and optimizing sequentially from the highest to the lowest priority. Both approaches target specific solutions along the Pareto front. In cases where the objectives are convex, the sequential optimization in (ii) can be reformulated as a weighted optimization in (i) with appropriate weightings [22]. In this section, we focus on the weighted method and demonstrate how to integrate the proposed two-stage mechanism with it.

Given a weighting vector $w = (w_1, \cdots, w_N) \succeq 0$, the element w_i represents the relative importance of the agent i's cost function. Without loss of generality, we let $\sum_{i=1}^N w_i = 1$. With the pre-assigned weightings, the planner aims to select the reference trajectories that minimize the weighted sum of agents' cost functions. This scenario is summarized in Mechanism 2, where the notation is carried over from Mechanism 1.

Theorem 2. *Mechanism 2 is strategyproof.*

Proof. The minimum cost player i can achieve by changing its decision y_i is

$$\min_{(y_i, \tilde{y}_{-i}^*(\tilde{c},\tilde{g})) \in \mathcal{F}_i} c_i\left(y_i, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) + p_i\left(y_i, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) \tag{13}$$

$$\geq \min_{y \in \mathcal{F}_i} c_i(y) + \sum_{j \neq i} \frac{w_j}{w_i} \tilde{c}_j(y) - h_i(\tilde{c}_{-i}, \tilde{g}_{-i}), \tag{14}$$

$$= c_i\left(\tilde{y}^*(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})\right) + p_i\left(\tilde{y}^*(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})\right), \tag{15}$$

where $\mathcal{F}_i = \{y : g_i(y) \preceq 0, \tilde{g}_j(y) \preceq 0, \forall j \neq i\}$ and $\tilde{y}^*(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})$ is the reference signal when agent i reports truthfully and (15) is the total cost of agent i. Note that (14) is just a scaling of (16) when agent i reports truthfully, therefore $\tilde{y}^*(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})$ is a minimizer of (14). As a result, reporting $\tilde{c}_i = c_i$ and $\tilde{g}_i = g_i$ is a weakly dominant strategy. ∎

A more complex scenario arises when the weightings are determined based on the reported functions, denoted as $w(\tilde{c}, \tilde{g})$. For example, if an unmanned aerial vehicle reports a low battery, it should receive a higher priority weighting. This type of scenarios is outlined in Mechanism 3. When weightings are not pre-assigned, agents may misreport to manipulate the planner's decision both directly, by distorting the optimization problem, and indirectly, by securing higher priority, which further influences the decision-making process. While VCG payments can resolve the issue of direct manipulation, addressing the indirect influence through priority requires additional measures. Before demonstrating our solution, we state two assumptions.

Assumption 3. $w_i(\tilde{c}, \tilde{g}) > \epsilon > 0$ *for all* i *with some fixed* ϵ *independent of* $\tilde{c}$ *and* $\tilde{g}$*. Equivalently, the ratios between all pairs of weightings are bounded by a constant* R*, i.e.,* $\frac{w_i(\tilde{c},\tilde{g})}{w_j(\tilde{c},\tilde{g})} \leq R, \forall i, j$.

Note that Assumption 3 ensures that each agent has a certain amount of importance and will not be neglected regardless of the information from all agents is reported.

Assumption 4. *If* $(\tilde{c}_i, \tilde{g}_i) \neq (c_i, g_i)$*, then* $c_i\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) + p_i\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) < c_i(\tilde{y}^*(\tilde{c}, \tilde{g})) + p_i(\tilde{y}^*(\tilde{c}, \tilde{g}))$*, for all* i.

This means that, if agent i reports cost/constraint functions untruthfully, the optimal decision made by agent i, y_i^*, will not match the reference suggested by the planner, $\tilde{y}_i^*(\tilde{c}, \tilde{g})$. A further discussion on Assumption 4 is given in the Appendix.

Theorem 5. *Suppose Assumption 3 and Assumption 4 hold. Then, Mechanism 3 is strategyproof.*

Proof. The minimum cost player i can achieve by changing

Mechanism 2 with pre-assigned weightings

(i) Each agent $i \in [N]$ reports its cost function $\tilde{c}_i$ and its constraint function $\tilde{g}_i$, which are not necessarily truthful. Then, based on the reported functions $\tilde{g} = (\tilde{g}_1, \cdots, \tilde{g}_N)$, $\tilde{c} = (\tilde{c}_1, \cdots, \tilde{c}_N)$, and the pre-assigned weighting w, the planner determines a reference signal

$$\tilde{y}^*(\tilde{c}, \tilde{g}) = \arg\min_{y} \sum_{i=1}^{N} w_i \tilde{c}_i(y) \quad s.t. \quad \tilde{g}(y) \preceq 0, \tag{16}$$

and broadcast it to the agents.

(ii) Upon receiving the reference signal from the planner, each agent i locally decides its optimal trajectory y_i^* by (6) and pays $p_i\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right)$ which is given by,

$$p_i(y_i, y_{-i}) = \sum_{j \neq i} \frac{w_j}{w_i} \tilde{c}_j(y_i, y_{-i}) + I_{\tilde{g}_j}(y_i, y_{-i}) - h_i(\tilde{c}_{-i}, \tilde{g}_{-i}), \tag{17}$$

where $h_i(\tilde{c}_{-i}, \tilde{g}_{-i})$ is an arbitrary function based on information reported by others, and

$$I_f(y_i, y_{-i}) = \begin{cases} 0 & \text{if } \ f(y_i, y_{-i}) \preceq 0 \\ \infty & \text{otherwise.} \end{cases}$$

Mechanism 3 with weightings based on reported information

(i) Each agent $i \in [N]$ reports its cost function $\tilde{c}_i$ and its constraint function $\tilde{g}_i$, which are not necessarily truthful. Then, based on the reported functions $\tilde{g} = (\tilde{g}_1, \cdots, \tilde{g}_N)$, $\tilde{c} = (\tilde{c}_1, \cdots, \tilde{c}_N)$, the planner decides a weighting $w(\tilde{c}, \tilde{g})$ and determines a reference signal

$$\tilde{y}^*(\tilde{c}, \tilde{g}) = \arg\min_{y} \sum_{i=1}^{N} w_i(\tilde{c}, \tilde{g}) \tilde{c}_i(y) \quad s.t. \quad \tilde{g}(y) \preceq 0, \tag{18}$$

and broadcast it to the agents.

(ii) Upon receiving the reference signal from the planner, each agent i locally decides its optimal trajectory y_i^* by (6) and pays $p_i\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right)$ which is given by,

$$p_i(y_i, y_{-i}) = \sum_{j \neq i} \frac{w_j(\tilde{c}, \tilde{g})}{w_i(\tilde{c}, \tilde{g})} \tilde{c}_j(y_i, y_{-i}) + I_{\tilde{g}_j}(y_i, y_{-i}) + \sum_{j \neq i} R \tilde{c}_j(y_i, y_{-i}) \delta(y_i, \tilde{y}_i^*(\tilde{c}, \tilde{g})) - h_i(\tilde{c}_{-i}, \tilde{g}_{-i}), \tag{19}$$

where $\delta(a, b) = 0$ if $a = b$ and 1 otherwise, $h_i(\tilde{c}_{-i}, \tilde{g}_{-i})$ is an arbitrary function based on information reported by others, and

$$I_f(y_i, y_{-i}) = \begin{cases} 0 & \text{if } \ f(y_i, y_{-i}) \preceq 0 \\ \infty & \text{otherwise.} \end{cases}$$

its decision y_i is

$$\min_{\left(y_i, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) \in \mathcal{F}_i} c_i\left(y_i, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) + p_i\left(y_i, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) \tag{20}$$

$$= c_i\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) - h_i(\tilde{c}_{-i}, \tilde{g}_{-i}) + \sum_{j \neq i} \left(\frac{w_j(\tilde{c}, \tilde{g})}{w_i(\tilde{c}, \tilde{g})} + R\delta\left(y_i^*, \tilde{y}_i^*(\tilde{c}, \tilde{g})\right) \right) \tilde{c}_j\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) \tag{21}$$

$$= c_i\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) - h_i(\tilde{c}_{-i}, \tilde{g}_{-i}) + \sum_{j \neq i} \left[\frac{w_j(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})}{w_i(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})} + A_{i,j} \right] \tilde{c}_j\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) \tag{22}$$

$$\overset{(a)}{\geq} c_i\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) - h_i(\tilde{c}_{-i}, \tilde{g}_{-i}) + \sum_{j \neq i} \frac{w_j(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})}{w_i(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})} \tilde{c}_j\left(y_i^*, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right) \tag{23}$$

$$\geq \min_{y \in \mathcal{F}_i} c_i(y) + \sum_{j \neq i} \frac{w_j(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})}{w_i(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})} \tilde{c}_j(y) - h_i(\tilde{c}_{-i}, \tilde{g}_{-i}) \tag{24}$$

$$= c_i\left(\tilde{y}^*(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})\right) + p_i\left(\tilde{y}^*(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})\right), \tag{25}$$

where $\mathcal{F}_i = \{y : g_i(y) \preceq 0, \tilde{g}_j(y) \preceq 0, \forall j \neq i\}$ and

$$A_{i,j} = \frac{w_j(\tilde{c}, \tilde{g})}{w_i(\tilde{c}, \tilde{g})} - \frac{w_j(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})}{w_i(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})} + R\delta\left(y_i^*, \tilde{y}_i^*(\tilde{c}, \tilde{g})\right).$$

To prove (a) holds, we prove $A_{i,j} \geq 0$ by considering the two cases: $y_i^* = \tilde{y}_i^*(\tilde{c}, \tilde{g})$ and $y_i^* \neq \tilde{y}_i^*(\tilde{c}, \tilde{g})$. When $y_i^* = \tilde{y}_i^*(\tilde{c}, \tilde{g})$, Assumption 4 implies that $\tilde{c}_i = c_i$ and $\tilde{g}_i = g_i$, leading to $A_{i,j} = 0$. When $y_i^* \neq \tilde{y}_i^*(\tilde{c}, \tilde{g})$, Assumption 3 ensures that $\frac{w_j(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})}{w_i(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})} \leq R$, resulting in $A_{i,j} \geq 0$. Note that $\tilde{y}^*(c_i, \tilde{c}_{-i}, g_i, \tilde{g}_{-i})$ is the reference signal when agent i reports truthfully and (25) is the total cost of agent i. As a result, reporting $\tilde{c}_i = c_i$ and $\tilde{g}_i = g_i$ is a weakly dominant strategy. ■

When agent i lies such that the weightings change, the importance of $\tilde{c}_j$ relative to c_i alters, but is upper bounded by R according to Assumption 3. The term $\sum_{j \neq i} R \tilde{c}_j\left(y_i, \tilde{y}_{-i}^*(\tilde{c}, \tilde{g})\right)$ penalizes the agent for indirectly influencing the planner's decision by affecting the determination of weightings.

V. Case Study: Constrained Linear-Quadratic-Regulator

We now demonstrate the proposed mechanisms using a MAS with linear dynamics, asymmetric coupling constraints, and a quadratic cost on a finite time horizon. The linear dynamics (4) are given as

$$x_i[k+1] = A_i x_i[k] + B_i u_i[k], \tag{26}$$

where $A_i \in \mathbb{R}^{4 \times 4}$, $B_i \in \mathbb{R}^{4 \times 2}$, $u_i[k] \in \mathbb{R}^2$ is the control input, and $x_i[k] = (p_i^x[k], p_i^y[k], v_i^x[k], v_i^y[k]) \in \mathbb{R}^4$ with its entries representing position and velocity in the x- and y-directions at time instant $k = 0, \cdots, K$. We assign each

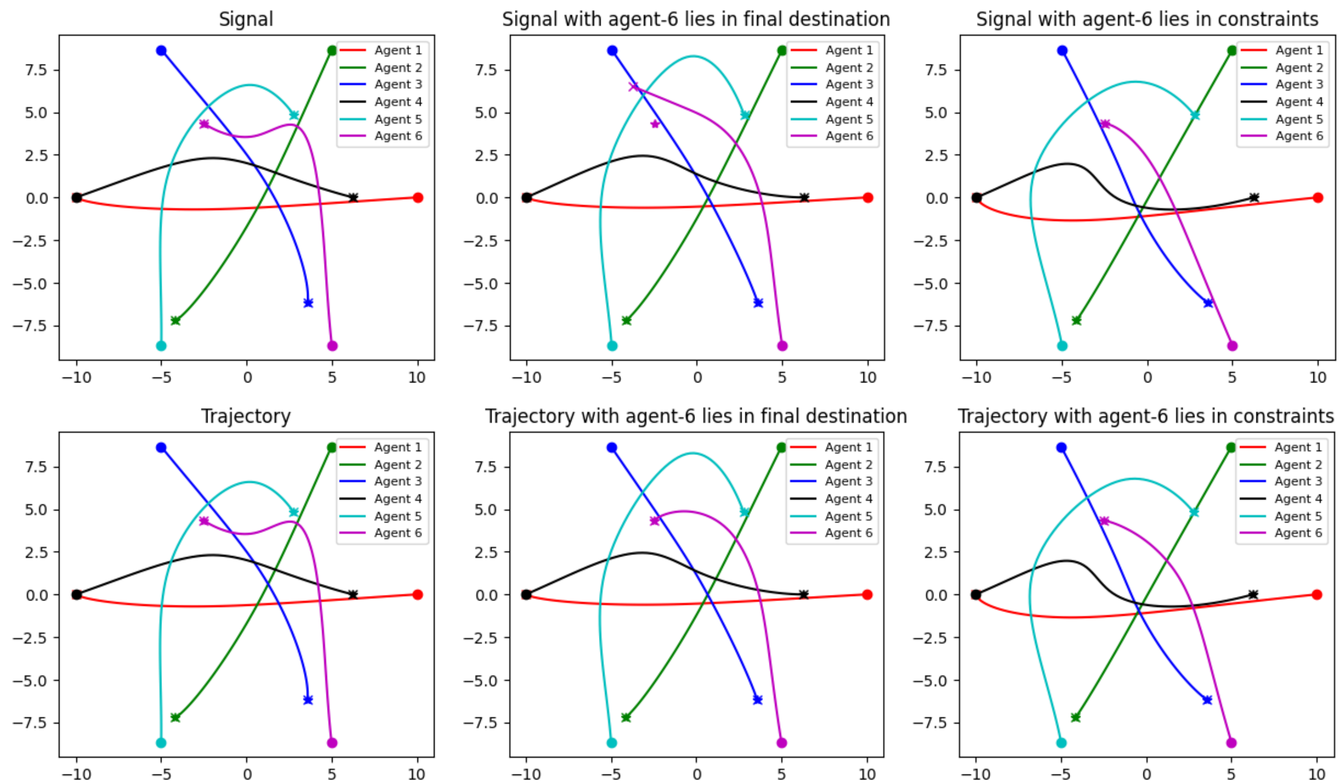

Fig. 1. This figure demonstrates when weightings are determined based on reported information. The upper row provides the signals generated from the planner while the lower row shows the trajectory optimized locally by the agents for the following three cases: (Left) All agents report truthfully, leading to the weightings $w = (32, 16, 8, 4, 2, 1)$. (Middle) Agent-6 lies in its final destination to get a higher priority, leading to the weightings $w = (32, 16, 8, 2, 1, 4)$. (Right) Agent-6 lies in its safe distance constraint to get a higher priority, leading to the weightings $w = (16, 8, 4, 2, 1, 32)$.

agent i a initial state x_i^0, a desired final state x_i^K and an input u_i^K that maintains it at that state, *i.e.*, $x_i^K = A_i x_i^K + B u_i^K$. The cost function $c_i(x, u)$ in (5) is quadratic:

$$c_i(x,u) = \sum_{k=0}^{K-1} \left(x_i[k] - x_i^K\right)^T Q_i \left(x_i[k] - x_i^K\right) + \left(u_i[k] - u_i^K\right)^T R_i \left(u_i[k] - u_i^K\right), \quad (27)$$

where $0 \preceq Q_i \in \mathbb{R}^{4\times 4}$ and $0 \prec R_i \in \mathbb{R}^{2\times 2}$. The coupling constraints $\bar{g}(x, u)$ in (5) encode a safe distance, d_i, from agent i to all others: $\forall i \neq j$ and $k = 0, \cdots, K$,

$$||p_i^x[k] - p_j^x[k]||^2 + ||p_i^y[k] - p_j^y[k]||^2 \geq d_i^2. \quad (28)$$

In addition, we impose hard constraints on confining agent i's final state to x_i^K. In the setting described above, the reported information from agent i to the planner is

$$\tilde{m}_i = \left(\tilde{A}_i, \tilde{B}_i, \tilde{Q}_i, \tilde{R}_i, \tilde{x}_i^0, \tilde{x}_i^K, \tilde{u}_i^K, \tilde{d}_i\right), \quad (29)$$

where the symbol "$\sim$" indicates that the information provided may not be truthful. The resulting optimization problems are nonlinear and nonconvex, and we solve them using IPOPT [23] and Couenne [24].

In Figure 1, we present the results when 6 weightings are determined based on the reported information and assigned by a permutation of the values $(1, 2, 4, 8, 16, 32)$. We suppose the agent with a farther destination or a smaller required safety distance is assigned a higher weighting. The upper row shows the signals generated by the planner, while the lower row displays the actual trajectories chosen by the agents.

When all agents report truthfully, the results are shown on the left side. In this case, the agent with a lower weighting follows a more curved path, as seen with agent 6. Therefore, agent 6 might misreport its destination to be farther to obtain a higher weighting. The results of this scenario are displayed in the middle, where agents 4 and 5 are forced to deviate from their preferred paths as agent 6 receives a higher weighting. Agent 6, however, ignores the signal it manipulated and optimizes its trajectory based on its true destination. Without additional penalties, agent 6 benefits from a lower-cost trajectory with misrepresentation.

One might expect the issue to be resolved using a VCG payment, which accounts for externalities. However, the result shows that even with a VCG payment, agent 6 can reduce its cost from 62.9 (truthful reporting) to 15.7 (misreporting). This outcome is caused by the agent influencing the weighting determination, which indirectly impacts the planner's decision. By applying Theorem 5 and imposing additional payments for manipulating the weightings, agent 6's cost rises to 169.9 when lying, making truthful reporting a better strategy.

In the scenario on the right, agent 6 falsely reports a lower safety distance to receive the highest weighting. Compared to the left case, all the other agents' signals are affected and agent 6 obtains a significantly better trajectory. As in the middle case, applying only the VCG payment still allows agent 6 to reduce its cost from 62.9 (truthful reporting) to 1.7 (misreporting). However, by applying the payment scheme from Theorem 5, agent 6 incurs a cost of 162.3 when misreporting, which is higher than reporting truthfully. It is important to note that in all three cases, agents other than agent 6 report truthfully. Therefore, the signals provide optimal solutions for these agents, as confirmed by the consistency between the executed trajectories (lower row) and the generated signals (upper row).

VI. Conclusions

We proposed a two-stage mechanism for MAS trajectory planning where agents have local decision-making power. Rather than strictly adhering to the centrally planned trajectory, agents can treat it as a recommendation to be refined, allowing them to adjust to local changes. We first proved that the proposed mechanism is strategyproof. Next, to accommodate the prevalent use of prioritized planning in the literature, we extended this mechanism to prevent the manipulation of priority weightings. In future research we will address *explicit* priority orders – rather than weightings – and explore mechanisms to prevent their manipulation.

References

[1] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[2] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.

[3] M. Kelly, "An introduction to trajectory optimization: How to do your own direct collocation," *SIAM Review*, vol. 59, no. 4, pp. 849–904, 2017.

[4] D. Malyuta, T. P. Reynolds, M. Szmuk, T. Lew, R. Bonalli, M. Pavone, and B. Açıkmeşe, "Convex optimization for trajectory generation: A tutorial on generating dynamically feasible trajectories reliably and efficiently," *IEEE Control Systems Magazine*, vol. 42, no. 5, pp. 40–113, 2022.

[5] P. Yang, K. Tang, J. A. Lozano, and X. Cao, "Path planning for single unmanned aerial vehicle by separately evolving waypoints," *IEEE Transactions on Robotics*, vol. 31, no. 5, pp. 1130–1146, 2015.

[6] M. Janner, Y. Du, J. Tenenbaum, and S. Levine, "Planning with diffusion for flexible behavior synthesis," in *International Conference on Machine Learning*, 2022.

[7] J. Park, J. Kim, I. Jang, and H. J. Kim, "Efficient multi-agent trajectory planning with feasibility guarantee using relative Bernstein polynomial," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 434–440.

[8] A. Majumdar and R. Tedrake, "Funnel libraries for real-time robust feedback motion planning," *The International Journal of Robotics Research*, vol. 36, no. 8, pp. 947–982, 2017.

[9] S. Singh, M. Chen, S. L. Herbert, C. J. Tomlin, and M. Pavone, "Robust tracking with model mismatch for fast and safe planning: an SOS optimization approach," in *Algorithmic Foundations of Robotics XIII: Proceedings of the 13th Workshop on the Algorithmic Foundations of Robotics 13*. Springer, 2020, pp. 545–564.

[10] K. S. Schweidel, H. Yin, S. W. Smith, and M. Arcak, "Safe-by-design planner–tracker synthesis with a hierarchy of system models," *Annual Reviews in Control*, vol. 53, pp. 138–146, 2022.

[11] Ángel Madridano, A. Al-Kaff, D. Martín, and A. de la Escalera, "Trajectory planning for multi-robot systems: Methods and applications," *Expert Systems with Applications*, vol. 173, p. 114660, 2021.

[12] N. Gunantara, "A review of multi-objective optimization: Methods and its applications," *Cogent Engineering*, vol. 5, no. 1, p. 1502242, 2018.

[13] R. Chai, A. Savvaris, A. Tsourdos, Y. Xia, and S. Chai, "Solving multiobjective constrained trajectory optimization problem by an extended evolutionary algorithm," *IEEE Transactions on Cybernetics*, vol. 50, no. 4, pp. 1630–1643, 2020.

[14] H. Hu, G. Dragotto, Z. Zhang, K. Liang, B. Stellato, and J. F. Fisac, "Who plays first? optimizing the order of play in Stackelberg games with many robots," in *Proceedings of Robotics: Science and Systems*, 2024.

[15] T. Kavuncu, A. Yaraneri, and N. Mehr, "Potential iLQR: A potential-minimizing controller for planning multi-agent interactive trajectories," *Robotics: Science and Systems*, 2021.

[16] P. Hespanhol and A. Aswani, "Surrogate optimal control for strategic multi-agent systems," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 2802–2807.

[17] K. Ma and P. R. Kumar, "The strategic LQG system: A dynamic stochastic VCG framework for optimal coordination," in *2018 IEEE Conference on Decision and Control (CDC)*, 2018, pp. 5777–5782.

[18] A. Camisa, P. N. Köhler, M. A. Müller, G. Notarstefano, and F. Allgöwer, "A distributed optimization algorithm for Nash bargaining in multi-agent systems," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 2684–2689, 2020, 21st IFAC World Congress.

[19] X. Ma, Z. Jiao, Z. Wang, and D. Panagou, "3-d decentralized prioritized motion planning and coordination for high-density operations of micro aerial vehicles," *IEEE Transactions on Control Systems Technology*, vol. 26, no. 3, pp. 939–953, 2018.

[20] O. Shorinwa and M. Schwager, "Scalable collaborative manipulation with distributed trajectory planning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 9108–9115.

[21] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of finance*, vol. 16, no. 1, pp. 8–37, 1961.

[22] K. Miettinen, *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers, 1999.

[23] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, pp. 25–57, 2006.

[24] P. Belotti, J. Lee, L. Liberti, F. Margot, and A. Wächter, "Branching and bounds tightening techniques for non-convex MINLP," *Optimization Methods and Software*, vol. 24, no. 4-5, pp. 597–634, 2009.

[25] W. Rudin, *Functional Analysis*, ser. International series in pure and applied mathematics. McGraw-Hill, 1991.

Appendix

To further examine Assumption 4, we first state the following proposition.

Proposition 1. *Let K be a compact set, $C(K)$ be all continuous functions over K, and $X = (C(K), ||\cdot||_\infty)$ be the metric space with infinity norm. Suppose for a given $f_0 \in C(K)$, we have $y_0^* = \arg\min_{y\in K} f_0(y)$. Then, the set of objective functions that have the same minimizer, $\mathcal{S}_{f_0} := \{f \in C(K) : y_0^* = \arg\min_{y\in K} f(y)\}$, is a meager set in X in the Baire category sense [25].*

To see the connection to Assumption 4, let f_0 represent the objective function for the central planner when agent i reports untruthfully, $\tilde{c}_i \neq c_i$. Correspondingly, y_0^* is the recommended trajectory for agent i provided by the planner, $\tilde{y}_i^*(\tilde{c}, \tilde{g})$. Assumption 4 supposes that $c_i \notin \mathcal{S}_{f_0}$, rendering $y_i^* \neq \tilde{y}_i^*(\tilde{c}, \tilde{g})$. This is not a restrictive assumption, since $\mathcal{S}_{f_0}$ is a meager set by the proposition above.