

Distributed Decision-Making Algorithms with Multiple Manipulative Actors: A Feedback-Control Perspective

Kasra Koorehdavoudi, Sandip Roy, Mengran Xue

Abstract—Distributed decision-making in the presence of multiple manipulative actors is studied, in the context of a linear distributed-consensus algorithm which has been enhanced to feedback controls enacted by these actors. The main contribution of the work is to evaluate the interplay among the manipulative actors in deciding the asymptotic decisions reached by the network of decision-making agents. In particular, the dependence of the asymptotic opinions on the network's topology and the manipulative actors' control schemes is characterized. Also, an example is used to illustrate that interactions among the actors may impact the dynamics of the decision-making algorithm in sophisticated ways.

I. INTRODUCTION

Distributed decision-making algorithms are used in applications ranging from computer networking to infrastructure health monitoring and multi-processor clock synchronization [1], [2]. Further, these algorithms are often descriptive of interactive processes that occur in human groups, such as voting processes within an organization or price consensus for commodity markets [3]. As the world around us becomes increasingly cyber enabled and networked, distributed processes for decision-making are becoming increasingly common.

The growing prominence of distributed decision-making algorithms has fostered a research effort on consensus and opinion dynamics in networks in the controls community [4]–[9]. One theme in this research effort has been to understand the influence of selfish stakeholders or actors, which aim to manipulate the decision process via their local actions [10]–[12]. Manipulation of decision-making processes has been modeled by including stubborn agents, whose opinions are unresponsive to their neighbors' opinions. Manipulation of decision-making algorithms has also been represented in a more general way as feedback controls (including specifically pinning controls) [13], or alternatively as open-loop actuations which can be analyzed using controllability concepts [9], [14]–[16].

Decision-making in many settings involves multiple selfish actors, which seek to manipulate network opinions toward divergent goals. For instance, voting in human groups almost always involves participants who influence others toward different decisions or candidates, and consensus is rarely reached. Likewise, in market processes, multiple actors may naturally seek to manipulate prices toward different values.

All authors are with the School of Electrical Engineering and Computer Science at Washington State University. The authors acknowledge support from United States National Science Foundation grants 1545104 and 1635184.

Thus, there is naturally an interest in understanding distributed decision-making processes with multiple manipulative actors. Based on this motivation, a few recent studies have considered distributed decision-making algorithms with multiple stubborn agents that have different opinions [14]. In a related direction, some studies have also considered algorithms that lead to group consensus, with different sets of agents being guided to different decisions or consensus values [10], [17].

This study is also concerned with distributed decision-making algorithms involving multiple selfish actors with divergent goals. The main contribution of the work is to understand the *interactions* among the selfish actors, i.e. to understand how the activities of one such actor influence the ability of another actor to manipulate the decision. Specifically, a linear distributed-consensus algorithm defined on a network is considered. The selfish actors are modeled as enacting feedback controls at one or a small set of agents. Our aim is to understand how the control capabilities of the different manipulative actors interact in deciding the outcome of the decision-making algorithm, and hence how the actions of some manipulative actors modulate the capabilities of others. Our primary focus here is on characterizing the decisions reached asymptotically by the agents, in terms of the controls used by the manipulative actors and the graph topology of the distributed decision-making algorithm. Via an example, we also illustrate that manipulative actors can alter the settling properties of the network and produce sophisticated impacts on transfer functions seen by other actors.

The paper is organized as follows. The distributed-decision-making model with manipulative actors is introduced and motivated in Section II. The main focus of the study, developed in Section III, is to understand how interactions among the manipulative actors decide the asymptotic decision achieved by each agent. In Section IV, an example is developed which illustrates the manipulation of the asymptotic decisions, and also explores the possibility for complex dynamic interactions among the actors. *Due to space constraints, proofs and some details are excluded, see [18].*

II. PROBLEM FORMULATION

A distributed decision-making algorithm specifies rules by which a network of *agents* update their *opinions* via local interactions, typically with the goal of reaching a common opinion or enabling an action (e.g. voting for a

leader) after some time. A number of distributed decision-making algorithms have been proposed, which may operate in continuous or discrete time, track continuous or quantal states, and use stochastic or deterministic update rules [19], [20]. Here, we enhance a standard deterministic continuous-time algorithm for continuous-valued opinions (termed a distributed consensus algorithm), to capture manipulation by multiple actors.

Formally, a network with n agents, specified by the set $\mathcal{N} = \{1, \dots, n\}$, is considered. Each agent i has an opinion $x_i(t)$ which evolves in continuous time ($t \in \mathbb{R}^+$). The interactions among the agents are defined by a weighted digraph $\Gamma = (V, E : W)$, where the weights W are assumed to be positive; we assume throughout the article that Γ is strongly connected. Nominally, each agent is modeled as updating its opinion based on a weighted average of differences between its opinion and those of its graphical neighbors. Mathematically, the nominal dynamics takes the form $\dot{\mathbf{x}} = -L(\Gamma)\mathbf{x}$, where $\mathbf{x} = [x_1 \ \dots \ x_n]^T$, and $L(\Gamma)$ is the (asymmetric) Laplacian matrix associated with the directed graph Γ . Specifically, $L(\Gamma)$ is defined as follows: the off-diagonal entry L_{ij} is set equal to the negative of the weight from vertex j to vertex i in the graph Γ if there is an edge (and to zero otherwise); the diagonal entries are selected so that each row sums to zero. This nominal model for distributed decision-making has been very widely studied in the literature [4], [9], [21]. Under broad connectivity conditions, the opinions of all agents can be shown to converge to a common value \bar{x} , which is a linear combination of their initial opinions. The dependence of the asymptotic value on the initial opinions can be prescribed through appropriate design of the network weights [4], [9].

Many decision-making processes have stakeholders that seek to manipulate the agents' opinion dynamics, to achieve decisions that differ from the one prescribed by a nominal algorithm. Here, we consider the circumstance that there are m independently-acting manipulative *actors*, specified by the set $\mathcal{M} = \{1, \dots, m\}$. Broadly, we view each actor as being able to access (actuate and measure) subsets or projections of the agents' opinions. Each actor is then modeled as enacting feedback control on a multi-input multi-output channel, with the aim of altering the steady-state or transient behaviors of the network in a way that provides a benefit relative to the other actors. In this study, we focus on one specific model for manipulative actors [9], [21], which are meant to capture typical type of manipulative behavior in distributed decision-making algorithms. The following is the model:

Local Actors: Each actor is able to influence a subset of the agents, using local feedback at each (i.e., the agent's feedback input depends only on its own opinion and the reference signal). Formally, in a distributed decision-making algorithm with local actors, each actor j ($j \in \mathcal{M}$) seeks to manipulate the agents in the set \mathcal{N}_j where $j \in \mathcal{M}$, to a common fixed point (constant reference) $r_j = \bar{r}_j$. The actor j applies a static (proportional) controller with gain k_{ij} at agent i . The agents' opinion dynamics with described local

actors is given by:

$$\begin{aligned} \dot{\mathbf{x}} &= -L(\Gamma)\mathbf{x} + \sum_{i \in \mathcal{N}} \mathbf{e}_i u_i \\ u_i &= \sum_{j \in \mathcal{M}} k_{ij}(\bar{r}_j - \mathbf{e}_i' \mathbf{x}), \end{aligned} \quad (1)$$

where k_{ij} is the specified gain if $i \in \mathcal{N}_j$ and k_{ij} is defined as 0 otherwise, and \mathbf{e}_i is 0–1 indicator vector with i th entry equal to 1. The distributed decision-making algorithm with local actors captures the circumstance that manipulators influence individual agents toward their preferred opinion, based on an understanding of their local opinion. We notice that the local-actor model is closely aligned with the stubborn-agent models in the literature [4], in the sense that the opinions of agents influenced by each actor are held close to the reference value specified by the actor. We also notice that the local-actors model can naturally be generalized to represent dynamics in the actors' influences on agents (e.g., sluggishness in the feedback actuation), by replacing the feedback gain with a transfer function $K_{ij}(s)$. The model with dynamic feedback also falls in the broad family of models with manipulating actors described above. We briefly discuss the model with dynamic local actors in Section III.

Our main goal in this work is to analyze how multiple manipulative agents in the network interact. The control enacted by each actor modifies the evolution of the distributed decision-making algorithm, and hence modulates the ability of the second actor to manipulate the dynamics. Our aim is to analyze these dependencies, focusing particularly on their impact on the asymptotic opinions achieved by the agents (Section 3). Via an example, we also introduce the idea that actions of one actor alter the dynamical responses caused by the actions of other actors (Section 4).

III. ANALYSIS OF ASYMPTOTIC OPINIONS

Our aim in this section is to study how the interactions among the manipulative actors decide the asymptotic behavior achieved by each agent. Nominally, if no manipulators were present, the agents would asymptotically achieve a common opinion or *decision*. The controls enacted by each actor modify the evolution of the distributed decision-making algorithm, and hence can modulate the asymptotic behaviors of the agents – perhaps preventing the formation of a steady state, or more typically leading to a steady-state with a gradation of opinions. Here, we seek to understand how the controls enacted by the different actors impact the asymptotic behavior of the agents, and thus to understand how one actors' efforts constrain the influence of the other actors. We present a sequence of analyses for the distributed decision-making algorithm with local actors.

In analyzing distributed decision-making algorithm with local actors (1), we assume that all local actors' control gains k_{ij} are nonnegative, and at least one of them is positive. This is the typical case that the local actors impose negative feedback controls, i.e. try to influence agents toward their reference signals. In this case, the state matrix of the controlled system dynamics is in the form of negative of

a grounded Laplacian matrix (associated with a strongly connected digraph). From properties of the grounded Laplacian matrix, it is immediate that the closed-loop system is asymptotically stable in the sense of Lyapunov. Since the reference inputs are constant signals, it follows immediately that the opinion of each agent reaches a steady-state. Further, from linearity, the steady-state opinion of each agent is seen to be a linear combination of the reference signals of all local manipulative actors. Hence, the asymptotic opinion \bar{x}_i of each agent (i.e., $\lim_{t \rightarrow \infty} x_i(t)$) can be presented as:

$$\bar{x}_i = \sum_{j \in \mathcal{M}} \lambda_{ij} \bar{r}_j. \quad (2)$$

Where \bar{x}_i is the asymptotic value for the opinion of the agent i ($x_i(t)$), i.e. $\lim_{t \rightarrow \infty} x_i(t) = \bar{x}_i$.

The following results characterize the weightings λ_{ij} in the asymptotic opinions of the agents in terms of the manipulating (local) actors' control gains and the network's graph topology. Hence, they give insight into the intertwined impacts of the manipulative actors on the agents' asymptotic opinions. To present these results, it is helpful to define some further terms: 1) $-L_{ij}$ is referred to as the *link-weight* of agent j on agent i , and 2) λ_{ij} is called the *contribution* of actor j to the asymptotic opinion of agent i . We notice that the link weight from j to i is equal to the network graph's edge weight if there is an edge from j to i , and is zero otherwise.

An initial result, developed in the following Theorem 1, is that the asymptotic opinions are weighted averages of the manipulative actors' reference signals:

Theorem 1: Consider the distributed-decision-making algorithm with local actors (1). For this model, consider the agents' asymptotic opinions, as defined in (2). The contribution of the local actor j to the asymptotic opinion of agent i , i.e. λ_{ij} , satisfies the following conditions: 1) $0 \leq \lambda_{ij} \leq 1$ and 2) $\sum_{j \in \mathcal{M}} \lambda_{ij} = 1$ for $\forall i \in \mathcal{N}$. In the case where each actor uses at least one strictly positive gain k_{ij} , the contributions λ_{ij} are strictly positive.

Remark: Theorem 1 has a close connection to the analysis in [14], which considers distributed decision-making algorithms with multiple stubborn agents. In particular, the stubborn-agent dynamics can be equivalenced with a static feedback control, albeit in a discrete-time rather than continuous-time setting. Hence a similar result on the dependence of the asymptotic opinions on the stubborn agents' initial opinions is recovered.

Theorem 1 has established that the asymptotic opinion of each agent in the network is a nonnegative unitary linear combination of the reference signals of the local actors. To understand the impacts of the manipulative actors on the asymptotic opinions, we next study the effect of the control gain of the local actors on the contributions λ_{ij} . The following theorem describes the dependence:

Theorem 2: Consider the distributed-decision-making algorithm with local actors (1). For this model, consider the agents' asymptotic opinions, as defined in (2). Suppose that there is more than one local actor where each actor uses

at least one strictly positive gain k_{ij} . The contribution of the local actor j to the asymptotic opinion of agent i , i.e. λ_{ij} , is a concave strictly-increasing function of the local actor's control gains k_{pj} for all positive k_{pj} , $p \in \mathcal{N}_j$ (i.e. $\frac{\partial \lambda_{ij}}{\partial k_{pj}} > 0$ and $\frac{\partial^2 \lambda_{ij}}{\partial k_{pj}^2} < 0$ for $\forall i \in \mathcal{N}$, $p \in \mathcal{N}_j$, and all positive k_{pj}). Meanwhile, each contribution λ_{ij} is a convex strictly-decreasing function of the other local actors' control gains k_{pq} for $q \neq j$, $p \in \mathcal{N}_q$ (i.e. $\frac{\partial \lambda_{ij}}{\partial k_{pq}} < 0$ and $\frac{\partial^2 \lambda_{ij}}{\partial k_{pq}^2} > 0$ for $q \neq j$, $\forall i \in \mathcal{N}$, $p \in \mathcal{N}_q$, and all positive k_{pq}).

The above theorem shows that, if a local actor q increases any of its control gains k_{pq} , its contribution to all agents' asymptotic opinions (i.e. λ_{iq} for $\forall i \in \mathcal{N}$) will increase or remain the same. Meanwhile, all other actors' contributions (i.e. λ_{ij} for $\forall j \neq q$ and $i \in \mathcal{N}$) will decrease or remain the same. Thus, by increasing a control gain, an actor increases his own influence on the agents and uniformly decreases the contributions of all other actors. It is interesting to ask whether an actor can gain absolute influence (i.e., achieve a contribution λ_{ij} of 1), by using a high gain. The following simple theorem clarifies that local actors can gain absolute influence on the agents which they directly actuate and measure.

Theorem 3: Consider the distributed-decision-making algorithm with local actors (1). For this model, consider asymptotic opinions of the agents p that are actuated/measured by an actor q ($p \in \mathcal{N}_q$ and $q \in \mathcal{M}$), as defined in (2). For sufficiently large control gains k_{pq} , the contribution of local actor q on agent p is arbitrarily close to 1, i.e. $\lim_{k_{pq} \rightarrow \infty} \lambda_{pq} = 1$.

Although actors can gain absolute influence on agents that they directly actuate/measure, their influence on other agents in the network is generally not absolute no matter how large a control gain they use. Thus, for agents that are not directly being actuated by an actor, other actors can ensure a minimum level of influence (minimum contribution to the asymptotic opinion of the actor) by applying a control anywhere in the network. Per Theorem 2, by increasing their control gains, these actors can also reduce the maximum influence achievable by the first actor.

The spatial pattern of impact of an actor, i.e. the dependence of the actor's contribution levels on the network's graph Γ , is also of interest. In the following theorem, we show the contribution of a local actor on the agents' asymptotic opinion is, in a certain sense, a monotonically decreasing function of the graphical distance of the agent from the actuation/measurement sphere of the actor. The theorem requires some terminology and notation relating to cutsets of the network's graph. To develop this terminology, let us consider a particular local actor j , and assume without loss of generality that the actor actually applies feedback to all agents that it actuates/measures (i.e. $k_{ij} \neq 0$ for $\forall i \in \mathcal{N}_j \subset \mathcal{N}$). Let us also define $\bar{\mathcal{N}}(j) \subset \mathcal{N}$ as the complement set of $\mathcal{N}(j)$. We define a *separating cutset* for actor j as a vertex cutset $\mathcal{Q}(j)$ such that all the vertices in subset $\mathcal{N}(j)$ are contained in only one partition formed by

the cutset, or on the cutset itself, i.e., all actuated vertices are on “one side” of the cutset or the cutset itself. (Notice that the vertices in subset $\mathcal{N}(j)$ may be present in all partitions formed by the cutset or on the cutset itself). We also refer to the corresponding group of the agents as *separating agents*. Let us use the label $\mathcal{S}(j)$ as the set of the vertices in the partition that includes vertices from set $\mathcal{N}(j)$, and refer to this set as the actor-close vertices (agents). Similarly, we use the label $\mathcal{T}(j)$ for the set of vertices in the partition(s) that do not include vertices from set $\mathcal{N}(j)$, and refer to this set as the actor-far vertices (agents). We note that any path from a vertex in set $\mathcal{S}(j)$ to a vertex in set $\mathcal{T}(j)$ passes through at least one of the vertices in set $\mathcal{Q}(j)$. Now the theorem is presented:

Theorem 4: Consider the distributed-decision-making algorithm with local actors (1). For this model, consider the agents’ asymptotic opinions, as defined in (2). Consider a particular local manipulative actor j and any corresponding separating cutset $\mathcal{Q}(j)$, which forms an actor-close set $\mathcal{S}(j)$ and an actor-far set $\mathcal{T}(j)$. The contribution of the local actor j to the asymptotic opinion of any agent in the actor-far set $\mathcal{T}(j)$ is less than its contribution to at least one separating agent specified in \mathcal{Q} , i.e. $\lambda_{ij} \leq \max_{\{q \in \mathcal{Q}(j)\}} \lambda_{qj}$ for $\forall i \in \mathcal{T}(j)$.

Theorem 4 shows that the influence of an actor is smaller at a remote agent, as compared to at least one agent on a cutset that separate the remote agent from its actuation/measurement set. This main result immediately implies that there is a spatial degradation in the contributions of actors, along sequential cutsets away from the agent’s measured/actuated set. The following three corollaries formalize this notion of a spatial degradation, beginning with a general case and then two specializations. The corollaries are presented without proof since they follow immediately from Theorem 4. The first corollary compares the contributions to the asymptotic opinion on two cutsets, in the general case:

Corollary 1: Consider the distributed-decision-making algorithm with local actors (1). For this model, consider the agents’ asymptotic opinions, as defined in (2). Consider an arbitrary local actor j and two separating cutsets $\mathcal{Q}_1(j)$ and $\mathcal{Q}_2(j)$, such that any path from vertices in set $\mathcal{N}(j)$, to the vertices in set $\mathcal{Q}_2(j)$ passes through at least one of the vertices in set $\mathcal{Q}_1(j)$. The contribution of local actor j to the asymptotic opinion of all separating agents in $\mathcal{Q}_2(j)$ is less than or equal to its contribution to at least one of the separating agents in $\mathcal{Q}_1(j)$, i.e. $\lambda_{ij} \leq \max_{\{q \in \mathcal{Q}_1(j)\}} \lambda_{qj}$ for $\forall i \in \mathcal{Q}_2(j)$.

The result can be presented more simply for the special case that the network graph Γ has a tree structure or (more generally) a single path between an actor’s measured/actuated agent and other agents of interest:

Corollary 2: Consider the distributed-decision-making algorithm with local actors (1). For this model, consider the agents’ asymptotic opinions, as defined in (2). Consider a particular local actor j and two agents p and q such that any path between a vertex in set $\mathcal{N}(j)$ to the vertex q passes through the vertex p . The contribution of the local actor j to

the asymptotic opinion of agent q is less than or equal to its contribution to the agent p , i.e. $\lambda_{qj} \leq \lambda_{pj}$.

Also, in case that each local actor is only actuating one agent, the following simplification can be obtained:

Corollary 3: Consider the distributed-decision-making algorithm with local actors (1). For this model, consider the agents’ asymptotic opinions, as defined in (2). Assume that local actor j only actuates the agent \hat{j} . The contribution of local actor j to the asymptotic opinion of its actuating agent \hat{j} , is greater than or equal to its contribution to the other agents’ asymptotic opinion, i.e. $\lambda_{ij} \geq \lambda_{qj}$ for all q and i .

The above results on the spatial pattern of impact for an actor (Theorem 4 and Corollaries 1-3) hold no matter the control gains used by the different actors. Hence, the results also hold for the maximum contribution that can be achieved by an actor on the agents through design of the local control gains.

In the previous analyses, we studied the effects of changes in a local manipulative actor’s control gain k_{ji} on the asymptotic opinions of the agents. Now, we investigate how changes in the link weight between two agents, i.e. entries of matrix $L(\Gamma)$, affect the asymptotic opinions in the network. These results give further intuition about how the network’s graph topology modulates the level of influence that different actors can achieve.

Theorem 5: Consider the distributed-decision-making algorithm with local actors (1). For this model, consider the agents’ asymptotic opinions, as defined in (2). The contribution of a local manipulative actor q on all agents’ asymptotic opinions is an increasing function of the link weight from agent i to agent j , i.e. $-L_{ji}$, if $\lambda_{iq} > \lambda_{jq}$ (i.e., $\frac{\partial \lambda_{uq}}{\partial L_{ji}} \geq 0$ in this case). Conversely, it is a decreasing function of $-L_{ji}$ if $\lambda_{iq} < \lambda_{jq}$.

Theorem 5 shows that the contribution of an actor to all agents’ asymptotic opinions (i.e., the overall influence of the actor) is increased if link weights from high-contribution agents to low-contribution agents are augmented. In other words, the actor’s influence can be spread widely if links are built from highly-influenced agents to other agents. Conversely, if less-influenced agents have stronger links to more influenced agents, then the influence of the actor globally decreases.

The next result specializes Theorem 5 to the case where each actor is only measuring/actuating one agent:

Corollary 4: Consider the distributed-decision-making algorithm with local actors (1). For this model, consider the agents’ asymptotic opinions, as defined in (2). Consider two actors \hat{i} and \hat{j} which each apply feedback controls at single agents, i and j respectively. The contribution of the actor i to all agents’ asymptotic opinions is a monotonically increasing function of $-L_{ji}$, while the contribution of actor \hat{j} to the asymptotic opinions is a monotonically decreasing function of $-L_{ji}$.

A further result is that, if the link weight from one agent to another is sufficiently increased, the asymptotic opinions of these two agents as well as the contributions of the local manipulative actors on these agents become close. The

following theorem states this result:

Theorem 6: Consider the distributed-decision-making algorithm with local actors (1). For this model, consider the agents' asymptotic opinions, as defined in (2). Consider any two agents i and j . The contribution of any local manipulative actor to the asymptotic opinion of the agents i and j , i.e. $\lambda_{i,q}$ and $\lambda_{j,q}$ for any $q \in \mathcal{M}$, become close to each other for sufficiently large $-L_{ij}$ or $-L_{ji}$. That is, $\lim_{(-L_{ij}) \rightarrow \infty} (\lambda_{iq} - \lambda_{jq}) = 0$ and $\lim_{(-L_{ji}) \rightarrow \infty} (\lambda_{iq} - \lambda_{jq}) = 0$ for any $q \in \mathcal{M}$.

Remark: Consider the distributed-decision-making algorithm and general local manipulating actors with feedback transfer functions $K_{ij}(s)$ as described in the problem formulation. Provided that the stability of the system with these actors is maintained, the asymptotic results presented in this section also easily follow through to the dynamic control case, with the DC control gain $K_{ij}(0)$ replacing the static gain.

IV. EXAMPLE AND PRELIMINARY DISCUSSION ABOUT DYNAMICS




Fig. 1. Graph Γ associated with a network with 7 agents and two actors.

An example is used to illustrate some results on the asymptotics of the distributed decision-making algorithm developed in Section III. The example is also used as a context for exploring, in a preliminary way, the dynamics of the distributed decision-making algorithm when manipulative actors are present. The example illustrates that manipulative actors may impact the global settling properties of the algorithm, as well as each others' control capabilities, in sophisticated ways.

For the example, a distributed-decision-making algorithm (1) defined on a network with 7 agents with two local manipulative actors is considered. The network's graph, which has a ring structure, is shown in Fig. 1: the weights in each direction are assumed to be identical, i.e. agents' influences are symmetric. Two local actors 1 and 2 are defined, which measure/actuate agents 1 and 2 respectively.

The contributions of the two actors to the agents' asymptotic opinions (i.e., λ_{ij}) have been determined, when the gains $k_{11} = 1$ and $k_{22} = 1$ are used: $\Lambda_{:,1} = [0.90, 0.09, 0.13, 0.15, 0.21, 0.26, 0.32]^T$, $\Lambda_{:,2} = [0.09, 0.90, 0.86, 0.84, 0.78, 0.73, 0.67]^T$. The contributions show that the influence of each actor on the agents' asymptotic opinions is decreasing with the distance from the agent actuated by the actor, as expected per Theorem 4. Additionally, the contribution of actor 2 to agent 5 is plotted as a function of the feedback gain k_{22} , for several different values of k_{11} , see Fig. 2. The figure shows that the actor's

influence on the agent is amplified as the feedback gain is increased, however the actor does not achieve absolute influence even with a large gain.




Fig. 2. The dependence of the contribution of local actor 2 on the asymptotic opinion of agent 5 on difference values of the control gains of local actor 1 and 2, i.e. k_{11} and k_{22} .

With regard to the dynamics of the distributed decision-making algorithm, the presence of manipulative actors may alter both the global dynamical properties of the algorithm (e.g. settling rate), and each others' control channels. Let us first use the example to explore the settling behavior. In the nominal model (without manipulative actors present), the settling rate is defined by the inverse of the subdominant eigenvalue of the graph Laplacian matrix, i.e. the non-zero eigenvalue which has smallest real part. For the example in Fig. 1, this subdominant eigenvalue is 0.1613. When manipulative actors are included in the model, the distributed decision-making algorithm (1) becomes asymptotically stable, and hence the settling rate is defined by inverse of the dominant eigenvalue of the grounded-Laplacian matrix associated with the model (the eigenvalue with smallest real part). For the example, the dominant eigenvalue with a single actor (Actor 1, using a gain $k_{11} = 1$) is 0.0206, while the dominant eigenvalue with both actors included ($k_{11} = 1$, $k_{22} = 1$) is 0.1079. Thus, we see that the presence of a single manipulative actor slows down the settling of the distributed decision-making algorithm, while the inclusion of the second manipulative actor comparatively speeds up the settling. While the rates are specific to the example, it can be shown via the Cauchy eigenvalue-interlacing theorem that the settling rate with a single local actor acting at a single agent will be slower than the nominal model, if the network graph is symmetric.

Beyond influencing the global dynamics of the distributed decision-making algorithm, the actions of one actor can also modulate the dynamics incurred by another actor's control in sophisticated ways. A natural way to study these dynamic interactions among actors is to characterize the impact of one actor's controllers on another actor's control channel, i.e. on the transfer function as seen from the perspective of the second actor. The poles of the transfer function are global properties which are intrinsic to the full network, while the transfer-function zeros are specific to the actor's control channel. Thus, understanding dynamic interactions among control channels requires characterizing the zeros of an actor's control channel in terms of the controls applied by

another actor. To illustrate the possible sophisticated impacts of one actor on another actor's control channel, we consider a slight modification of the example above. In particular, we also include a new actor (Actor 3) which measures/targets agent 3 and actuates agent 7. We study how the zeros of the transfer function for actor 3 vary with the control gain k_{22} used by actor 2 (while the control gain used by actor 1 is fixed at $k_{11} = 1$). In particular, the dependence of the dominant zero location (the largest real part among the zeros) on k_{22} is shown in Fig. 3. The plot shows that the transfer function seen by actor 3 is nonminimum phase for small gains k_{22} (i.e. has zeros in the right half plane), but becomes minimum-phase for large gains. It is well-known that the presence of nonminimum-phase zeros complicates control: for instance, leading to instability if static high-gain controllers are used, and placing limits on control performance (e.g. on reference tracking, disturbance rejection, or cheap control). Thus, in this example, actor 2 can support actor 3's efforts to manipulate the network by using a higher gain. More generally, the example highlights that multiple actors may have sophisticated dynamic interactions, which may: 1) lead to unexpected responses to an actor's feedback control, 2) limit or amplify actors' opportunities for manipulation, and 3) require changed control topologies of the actors. A comprehensive analysis of such dynamic interactions is left to future work: we refer the reader to our earlier work on control-channel interactions for some results that are relevant to the dynamic analysis [21]–[26].




Fig. 3. The dependence of the dominant zero location (the largest real part among the zeros) for transfer function channel 1-3 on the control gain k_{22} of actor 2.

REFERENCES

- [1] N. A. Lynch, *Distributed algorithms*. Elsevier, 1996.
- [2] G. A. Pratt and J. Nguyen, "Distributed synchronous clocking," *IEEE transactions on parallel and distributed systems*, vol. 6, no. 3, pp. 314–328, 1995.
- [3] R. A. Holley and T. M. Liggett, "Ergodic theorems for weakly interacting infinite systems and the voter model," *The annals of probability*, pp. 643–663, 1975.
- [4] A. V. Proskurnikov and R. Tempo, "A tutorial on modeling and analysis of dynamic social networks. part i," *Annual Reviews in Control*, vol. 43, pp. 65–79, 2017.
- [5] V. D. Blondel, J. M. Hendrickx, A. Olshevsky, and J. N. Tsitsiklis, "Convergence in multiagent coordination, consensus, and flocking," in *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC'05. 44th IEEE Conference on*. IEEE, 2005, pp. 2996–3000.
- [6] W. Ren and R. W. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Transactions on automatic control*, vol. 50, no. 5, pp. 655–661, 2005.

- [7] V. D. Blondel, J. M. Hendrickx, and J. N. Tsitsiklis, "Continuous-time average-preserving opinion dynamics with opinion-dependent communications," *SIAM Journal on Control and Optimization*, vol. 48, no. 8, pp. 5214–5240, 2010.
- [8] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.
- [9] M. Mesbahi and M. Egerstedt, *Graph theoretic methods in multiagent networks*. Princeton University Press, 2010.
- [10] S. Sundaram and C. N. Hadjicostis, "Distributed function calculation via linear iterative strategies in the presence of malicious agents," *IEEE Transactions on Automatic Control*, vol. 56, no. 7, pp. 1495–1508, 2011.
- [11] M. Pirani and S. Sundaram, "Spectral properties of the grounded laplacian matrix with applications to consensus in the presence of stubborn agents," in *American Control Conference (ACC), 2014*. IEEE, 2014, pp. 2160–2165.
- [12] J. Ghaderi and R. Srikant, "Opinion dynamics in social networks: A local interaction game with stubborn agents," in *American Control Conference (ACC), 2013*. IEEE, 2013, pp. 1982–1987.
- [13] F. Chen, Z. Chen, L. Xiang, Z. Liu, and Z. Yuan, "Reaching a consensus via pinning control," *Automatica*, vol. 45, no. 5, pp. 1215–1220, 2009.
- [14] J. Ghaderi and R. Srikant, "Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate," *Automatica*, vol. 50, no. 12, pp. 3209–3215, 2014.
- [15] R. Dhal and S. Roy, "Vulnerability of continuous-time network synchronization processes: A minimum energy perspective," in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*. IEEE, 2013, pp. 823–828.
- [16] F. Pasqualetti, S. Zampieri, and F. Bullo, "Controllability metrics, limitations and algorithms for complex networks," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 40–52, 2014.
- [17] J. Yu and L. Wang, "Group consensus in multi-agent systems with switching topologies and communication delays," *Systems & Control Letters*, vol. 59, no. 6, pp. 340–348, 2010.
- [18] K. Koorehdavoudi, S. Roy, M. Xue, and J. A. Torres, "Distributed decision-making algorithms with multiple manipulative actors," in *preparation*.
- [19] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on automatic control*, vol. 49, no. 9, pp. 1520–1533, 2004.
- [20] S. Roy, K. Herlugson, and A. Saberi, "A control-theoretic approach to distributed discrete-valued decision-making in networks of sensing agents," *IEEE Transactions on Mobile Computing*, vol. 5, no. 8, pp. 945–957, 2006.
- [21] K. Koorehdavoudi, S. Roy, J. A. Torres, and M. Xue, "Interactions among control channels in dynamical networks," in *Decision and Control (CDC), 2017 IEEE 56th Annual Conference on*. IEEE, 2017, pp. 1058–1063.
- [22] K. Koorehdavoudi, M. Hatami, S. Roy, V. Venkatasubramanian, P. Panciatici, F. Xavier, and J. A. Torres, "Input-output characteristics of the power transmission network's swing dynamics," in *Decision and Control (CDC), 2016 IEEE 55th Conference on*. IEEE, 2016, pp. 1846–1852.
- [23] K. Koorehdavoudi, S. Roy, T. Prevost, F. Xavier, P. Panciatici, and V. Venkatasubramanian, "Input-output properties of the swing dynamics for power transmission networks with hvdc modulation," in *Proceedings of the 20th IFAC World Congress*, 2017.
- [24] K. Koorehdavoudi, M. Yao, S. Roy, and J. L. Mathieu, "Using demand response to shape the fast dynamics of the bulk power network," in *Proceedings of the IREP Symposium on Bulk Power System Dynamics and Control*, Espinho, Portugal, 2017.
- [25] K. Koorehdavoudi, S. Roy, J. A. Torres, and M. Xue, "Influence of diffusive-network controllers on remote-channel zeros," in *preparation*.
- [26] —, "Impacts of high and low gain controllers on remote channels in dynamical networks," in *IEEE Conference on Decision and Control (to appear)*, 2018.