Collective Decision-Making in Multi-Agent Systems by Implicit Leadership

Chih-Han Yu
SEAS & Wyss Institute
Harvard University
chihanyu@gmail.com

Justin Werfel
Wyss Institute
Harvard University
justin.werfel@wyss.harvard.edu

Radhika Nagpal
SEAS & Wyss Institute
Harvard University
rad@eecs.harvard.edu

ABSTRACT

Coordination within decentralized agent groups frequently requires reaching global consensus, but typical hierarchical approaches to reaching such decisions can be complex, slow, and not fault-tolerant. By contrast, recent studies have shown that in decentralized animal groups, a few individuals without privileged roles can guide the entire group to collective consensus on matters like travel direction. Inspired by these findings, we propose an implicit leadership algorithm for distributed multi-agent systems, which we *prove* reliably allows all agents to agree on a decision that can be determined by one or a few better-informed agents, through purely local sensing and interaction. The approach generalizes work on distributed consensus to cases where agents have different confidence levels in their preferred states. We present cases where informed agents share a common goal or have conflicting goals, and show how the number of informed agents and their confidence levels affects the consensus process. We further present an extension that allows for fast decision-making in a rapidly changing environment. Finally, we show how the framework can be applied to a diverse variety of applications, including mobile robot exploration, sensor network clock synchronization, and shape formation in modular robots.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Multiagent systems

General Terms

Algorithms, Performance, Theory

Keywords

Biologically-inspired approaches and methods, Multi-robot systems, Collective Intelligence, Distributed Problem Solving

1. INTRODUCTION

Animals achieve robust and scalable group behavior through the distributed actions of many independent agents, e.g., bird flocking, fish schooling, and firefly synchronization. In such systems, each agent acts autonomously and interacts only with local neighbors, while the global system exhibits coordinated behavior and is highly adaptive to local changes. It is a remarkable empirical fact that the actions of a large group can be determined by one or a few "informed" individuals with privileged knowledge, without explicit signaling [2]. Here we present an algorithm that exploits this type of decision-reaching process, which we call implicit leadership, and *prove* that a system of agents following it will come to rapid global agreement under appropriate assumptions. This work also generalizes the formalism of distributed consensus [3, 9, 16], a method shown to be successful in multi-agent agreement tasks, and extends it to handle cases where agents may have different degrees of confidence in their preferred state.

Coordination can be challenging in large multi-agent systems where agents are spatially distributed or heterogeneous, as some agents will be able to obtain important information not available to others. For example, in a multi-robot search task, robots close to a target may be able to detect or track it better than those farther away; in sensor network time synchronization tasks, nodes closer to the base station can provide more precise timing information. These *informed* agents need to play a more important role when it comes to reaching a group consensus. One approach that can incorporate this privileged information is to designate a centralized coordinator to collect information from all agents, then disseminate a decision to the whole group. However, such a strategy interferes with the system's scalability and robustness: the coordinator can easily become a communication bottleneck, and it is also a potential point of failure for the system [15].

Natural systems can suggest decentralized control laws that avoid these drawbacks. Couzin et al. have proposed a model for animals in groups such as migrating herds, in which only a few informed individuals influence the whole group's direction of travel [2]. Ward et al. report on nonlinear decision mechanisms in which fish choose movement directions based on the actions of a quorum of their neighbors, some of whom may have privileged information [14]. These studies provide examples of effective decentralized decision-making under information asymmetry, in which there is *no* explicit announcement of leaders or elaborate communication; each individual simply modifies its behavior based on that of its neighbors.

One disadvantage of these studies from an engineering standpoint is that they are purely empirical, without guarantees regarding the observed behavior or analytic treatment of the conditions under which the results will hold. Designing effective controllers for artificial systems requires addressing several open questions: How do we prove that the informed agents correctly influence the rest of the group? Can we characterize factors related to the speed with which the group can reach a decision? How do we resolve cases involving conflicting goals?

Here we propose an implicit leadership framework by which distributed agents can reach a group decision efficiently. We address the above theoretical challenges and generalize the approach to more complex scenarios. Our approach is that of distributed consensus (DC) algorithms [3, 9, 16], with a critical departure: while

Cite as: Collective Decision-Making in Multi-Agent Systems by Implicit Leadership, Chih-Han Yu, Justin Werfel, and Radhika Nagpal, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.

traditional DC studies make the assumption that all agents have equally valid information, here we allow agents to be informed to varying degrees. This generalization allows DC approaches to be applied to a wider range of scenarios. We characterize the algorithm's theoretical properties and provide extensions: (1) We prove that its use indeed leads to the desired group behavior: a subset of informed agents with the same goal can correctly lead the whole group to adopt that goal, achieving a consensus decision. (2) We characterize how factors such as the fraction of informed agents in the group affect the speed with which the group arrives at a decision. (3) We generalize the algorithm to the case where informed agents have different goals, by having them modify their confidence levels according to the states of their neighbors (*implicit leadership reinforcement*). (4) We further extend the framework to improve group performance and convergence speed in rapidly changing environments (*fast decision propagation*). Finally, we demonstrate how the framework can be applied to a diverse variety of distributed systems applications, including multi-robot exploration, sensor network clock synchronization, and adaptive shape formation in modular robots.

2. RELATED WORK

Biological collective behaviors, where coordinated global behavior emerges from local interactions, have inspired decentralized approaches in various types of multiagent systems, e.g., swarm robotics [1,7], modular robots [12,16], and sensor networks [6]. A central challenge for these systems is designing local agent control rules that will *provably* lead to desired global behaviors.

In the control theory community, there has been some success in analyzing bio-inspired rules related to distributed consensus, e.g., in flocking [8] and multi-robot formation tasks [3, 9]. Rules proposed for these problems have similar algorithmic properties and their convergence and performance have been analyzed for arbitrary topologies. While most systems assume leaderless scenarios, some recent studies have shown how to control the whole group's behavior by controlling specific agents designated as leaders, whether a single agent [3], multiple agents [4,13], or even virtual leaders not corresponding to actual agents [5]. Such studies generally assume that leaders are pre-designated and have the same objectives. In a distributed multi-agent system, agents that can sense important aspects of the environment might naturally provide guidance to the group, without an explicitly privileged role as leaders. Inspired by recent biological studies [2,14], we propose and analyze an implicit leadership framework that can capture such scenarios, and further show how these mechanisms can resolve conflicts between leaders and achieve fast decision-making.

Particle Swarm Optimization (PSO) shares certain similarities with our algorithms in terms of mathematical formulation [10]. However, our overall framework is quite distinct from it: PSO is an optimization technique with each particle (agent) representing a solution to the objective function, while our framework addresses distributed multiagent decision-making tasks with each agent representing an independent decision maker.

3. MULTI-AGENT MODEL

Our model comprises a set of independent agents a_i, each of which can control certain of its state variables x_i (e.g., heading direction in Fig. 1), and determine the states of neighbors within some fixed range. This determination may be passive, via sensing, or active, via communication. Some agents have privileged knowledge, e.g., due to better sensing capabilities or a location close to an object of interest. These informed agents have a goal state x_i^*,

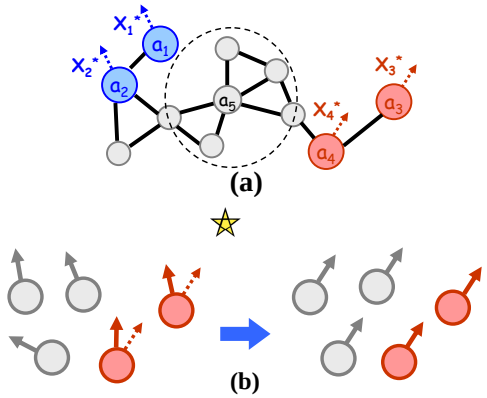

Figure 1: (a) A node indicates an agent and an edge indicates a neighborhood link. The dotted line indicates agent a_5's sensor/communication range. Informed agents (colored) have their own goal states (arrows). (b) Agents start with random headings and try to make a consensus decision. Solid/dotted arrows for informed agents represent current/goal directions, respectively.

Algorithm 1 Equation for state update with implicit leadership. All agents execute this update at each time step. Informed agents ($w_i > 0$) influence the rest of the group without explicit commands. The two terms represent behaviors (a) and (b) as described in the text.

loop

$$x_i(t+1) \leftarrow \underbrace{\frac{1}{1+w_i} \cdot x_{\text{avg}}(t)}_{\text{Behavior (a)}} + \underbrace{\frac{w_i}{1+w_i} \cdot x_i^*}_{\text{Behavior (b)}}$$

and a confidence w_i in that goal (Fig. 1 (a)).

Neighborhood relationships among agents can be described by a connectivity graph (Fig. 1 (a)). Nodes represent agents, and edges indicate that agents are within sensor or communication range. We assume that this relationship is symmetric; hence the graph is undirected. Since agents can move, the graph can potentially be time-varying. We refer to this case as a *dynamic topology*. If the graph remains static, we refer to it as a *static topology*.

Throughout the following sections, we use *collective pursuit* in multi-robot systems as a motivating example: there exists some moving target whose heading direction the whole group needs to match (Fig. 1 (b)). Such a task may arise in, e.g., patrol scenarios (finding and pursuing an intruder), exploration (where the group moves in the direction of a feature of interest), or human-robot interaction (where a user needs to guide a swarm to a desired location). In this example, an agent's (robot's) direction is its relevant state variable. In §7 we discuss applying our approach to a wide variety of tasks.

4. IMPLICIT LEADERSHIP ALGORITHM

Here we present the algorithm that achieves collective decision-making in the multiagent system. Inspired by Couzin et al. [2], our agent control law incorporates two components: (a) a tendency to align its state with those of its neighbors, and (for informed agents) (b) an attempt to achieve its goal state. At each time step, each agent computes its new state as a function of its old state and the states of its neighbors. The control algorithm can be formally written as in Algorithm 1. In that equation, a_i indicates agent i and $x_{\text{avg}}(t)$ is simply the average of the states of agent a_i and its neigh-

bors: $x_{\text{avg}}(t) = \left(x_i(t) + \sum_{j|a_j \in \mathcal{N}_i(t)} x_j(t)\right) / (|\mathcal{N}_i(t)| + 1)$, where $\mathcal{N}_i(t)$ is the set of a_i's neighbors at time step t. The factor $w_i/(1+w_i)$ captures the degree to which agent a_i *drives itself to the goal state* x_i^*, and $1/(1+w_i)$ the degree to which it *follows the local average*. If a_i is an informed agent, its confidence w_i is positive; otherwise $w_i = 0$. In the rest of this section, we assume that w_i and x_i^* are constant over time.

All agents execute the simple rule of Algorithm 1, with no need for informed agents to take specific actions to coordinate others or communicate their goal states. For instance, when a robot team starts with randomly initialized heading directions and only some robots have knowledge of a unique global goal direction, the latter set w_i to be positive to reflect their confidence and x_i^* to match that goal. All agents repeatedly update their heading direction $x_i(t)$ according to Algorithm 1, and are iteratively influenced by the informed agents to reach the desired consensus. Each agent knows it has reached a stable decision state when the maximal state difference between it and any of its neighbors is less than a small number ϵ.

4.1 Algorithm Analysis

Here we analytically prove properties of Algorithm 1's correctness and performance. In this section we state our theorems, and defer all proofs to the Appendix.

First, in the absence of informed agents ($w_i = 0\,\forall i$), our control formulation becomes identical to that of *distributed consensus* [3, 9, 16]. If the communication graph is connected or periodically-connected[1], all agents will converge to the same state: $\lim_{t\to\infty} x_i(t) = \bar{x} = \frac{1}{N}\sum_i x_i(0)\,\forall i$, where N is the number of agents. Distributed consensus analyses have been applied to many areas, e.g., sensor network synchronization, and the theoretical properties have been widely explored [9].

Next, we consider how the decision of the group is influenced when there are informed agents all with the same goal state ($x_i^* = x^*\,\forall i$ such that $w_i > 0$). More specifically, we ask: (1) Will all agents converge to x^*? (2) How does convergence speed relate to the confidence factors w_i and the number of informed agents? (3) Will agents converge to the same state in both dynamic and static topology cases?

We first introduce two definitions: (1) Informed agent group: A set of informed agents $\mathcal{L}_k$ with the same goal state $\bar{x}_k^*$. (2) Informed agent-connected: A condition for dynamic topologies in which a path always exists between each non-informed agent and some informed agent, and this path persists without modification for $d_{\max} + 2$ time steps. Here $d_{\max}$ is the maximum among all non-informed agents of the shortest hop distance to an informed agent.

THEOREM 4.1. *(Convergence) Let there be only one informed agent group with goal state x^*, with all agents executing control algorithm 1. Let the communication topology be either (1) static and connected, or (2) dynamic and informed agent-connected. Then*

$$\lim_{t\to\infty} x_i(t) = x^*\,\forall i$$

PROOF. See Appendix A. □

Intuitively, Theorem 4.1 tells us that if the agent topology is either connected or informed agent-connected, all agents will eventually be able to follow a single informed agent group.

Next we consider how the group size and confidence weights of informed agents can affect the convergence rate.

[1]A graph with dynamic topology is periodically-connected if, for all pairs of nodes $\{a_i, a_j\}$, some path P between them exists and persists without modification for a number of time steps $\geq$ length$(P) + 1$.

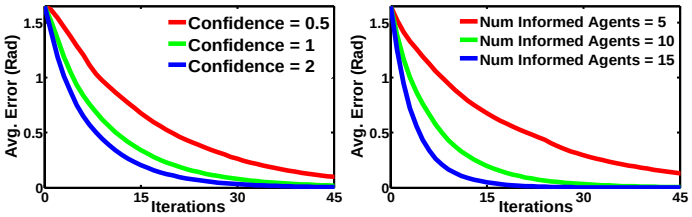

(a) (b)

Figure 2: Average error in heading direction within the group over time, for (a) different confidence levels and (b) number of informed agents. The y-axis gives agents' average distance from the goal state. Convergence speed increases when the confidence and/or number of informed agents increases. Each data point is averaged from 50 random initializations.

THEOREM 4.2. *Under the same conditions as Theorem 4.1, the system's convergence rate increases if more agents become informed, or if existing informed agents increase their confidence levels, in both connected (static) and informed agent-connected (dynamic) topologies.*

PROOF. See Appendix B. □

Experimental Results: We also explore these relationships in simulation. Our simulation model is composed of 50 agents that are initially randomly distributed within a 10×10 area (arbitrary units). They can freely travel in 2D space without boundaries. Each agent moves a distance of 0.05 per iteration and has a communication radius of 4. An agent's state is characterized by its heading direction $x_i \in [0, 2\pi)$. All agents begin in random states, and informed agents are given a common goal direction. Our evaluation metric is defined as agents' average distance from the goal state. Fig. 2 shows how the convergence speed changes when (a) all informed agents' confidence weights are increased (we fix the number of informed agents at 10), or (b) more agents are selected as informed agents (we fix informed agents' confidence at 2). The convergence speed to the informed agents' goal direction increases in both cases.

Finally we consider the case where informed agents do not all share the same goal state—i.e., multiple informed agent groups.

THEOREM 4.3. *(Multiple Informed Agent Groups) Let there be multiple informed agent groups $\mathcal{L}_1, \mathcal{L}_2, \cdots, \mathcal{L}_q$ with different goal states $\bar{x}_1^*, \bar{x}_2^*, \cdots, \bar{x}_q^*$. Let agent topology be static and connected. Then each agent's state will converge to some value, but this value can be different for different agents.*

PROOF. See Appendix C. □

Consider an example with agents arranged in a line, where the two agents on the ends are informed and have opposite goal directions. The agents between them will then converge to states in between the goal states of the two informed agents. This convergence holds for static topologies, but not for dynamic topologies; in the latter the agents may never converge to a fixed state[2].

Given these results, if a single group consensus is to be reached and there are initially multiple informed agent groups, it is necessary for them to become a single group with a common goal state.

[2]To see this, suppose that each agent's state has converged to a fixed value, and that the topology then changes such that a non-informed agent a_i goes from having an informed neighbor in $\mathcal{L}_j$ to instead having one in $\mathcal{L}_k (j \neq k)$. Then a_i's state will clearly change.

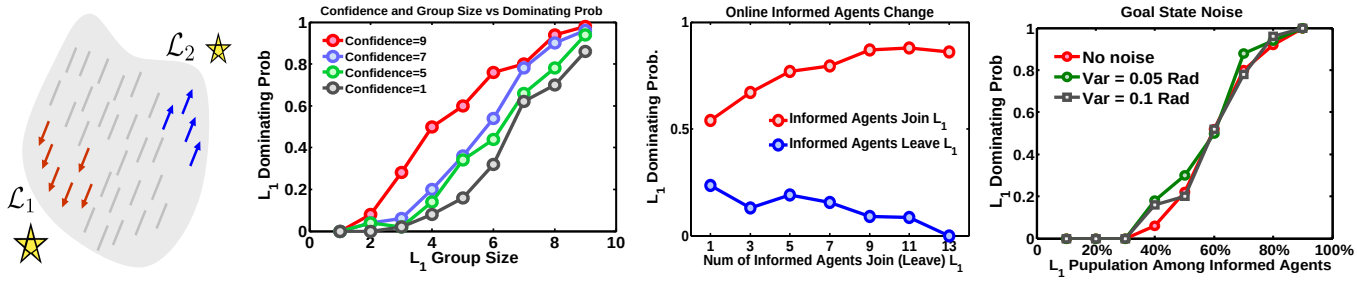

(a) (b) (c) (d)

Figure 3: (a) $\mathcal{L}_1$ and $\mathcal{L}_2$ observe different targets and have distinct goal states. (b) When informed agents in $\mathcal{L}_1$ increase their confidence (different curves) or the size of $\mathcal{L}_1$ increases (x-axis), it increases the probability that $\mathcal{L}_1$ dominates (i.e., ultimately becomes the only remaining informed agent group). (c) When more agents join $\mathcal{L}_1$ during the course of a run, $\mathcal{L}_1$'s dominating probability increases; similarly, the probability decreases when more agents leave $\mathcal{L}_1$. (d) The probability that $\mathcal{L}_1$ dominates is not appreciably affected by slight variation among the goal states of agents within the group.

Algorithm 2 Implicit leadership reinforcement (ILR) lets informed agents change their confidence so that the entire group can reach a single consensus. Agents execute this reinforcement routine along with the standard implicit leadership algorithm (Algorithm 1). The function $1\{C\}$ evaluates to 1 if condition C is true and 0 otherwise (line 3). We set $\delta = 0.6$, $\bar{\rho} = 0.4$, $\bar{t} = 5$, and $\alpha = 0.5$.

loop
  if a_i is an informed agent **then**
    if $\sum_{j|a_j \in \mathcal{N}_i} \frac{1\{\|x_j(t) - x_i^*\| \leq \delta\}}{|\mathcal{N}_i|} > \bar{\rho}$ **then**
      $w_i(t+1) = w_i(t) + \alpha$
    else
      $w_i(t+1) = \max(w_i(t) - \alpha, 0)$
    if w_i has been 0 for $\bar{t}$ time steps **then**
      a_i becomes non-informed ($w_i = 0$ permanently)

In the following, we describe how to extend Algorithm 1 to achieve this.

5. IMPLICIT LEADERSHIP REINFORCEMENT (ILR) ALGORITHM

In many tasks, it is necessary for a single informed agent group to emerge. For example, in a multi-robot search and rescue task, informed robots at different locations may observe different clues and thus have different opinions about where rescue subjects are located, yet the whole group may eventually need to agree on a single rescue target to secure enough robot force to complete its task. We can achieve this aim by replacing the constant w_i of informed agents with a time-varying variable $w_i(t)$. We call the mechanism that decides how $w_i(t)$ should change *implicit leadership reinforcement* (ILR).

Algorithm 2 describes ILR. The idea is that informed agents have their confidence increase if the states of a sufficiently large fraction $\bar{\rho}$ of their neighbors are sufficiently close to their own goal states (within δ); otherwise the confidence decreases, and if it reaches 0 for long enough then the agent loses its "informed" status.

This approach has the following advantages. (1) While each agent has only a local view, the result that emerges nevertheless captures *global statistics*. The informed agent group with a larger population or higher confidence dominates the decision with a high probability. (2) At any time, agents can join or leave an informed agent group or change confidence while observing more instances, and the system will accommodate these updates without the process needing to restart. (3) Only a few iterations are required for a single informed agent group to emerge.

We call an informed agent group a *dominating group* if at some point it contains all remaining informed agents—i.e., all agents in other groups have become non-informed. If this occurs, then all agent states will converge to the dominating group's goal state so long as the connection graph is connected or informed agent-connected (Theorem 4.1).

Experimental Results: Since we do not yet have analytical proofs of convergence of ILR, we investigate its performance experimentally. We use the same simulation setup as in Section 4.1. Initially, two informed agent groups $\mathcal{L}_1$ and $\mathcal{L}_2$ are constructed by choosing a total of ten informed agents forming two spatial clusters (Fig. 3 (a)). The two groups are given distinct goal headings, set in opposite directions. In 1800 experiments, the average number of iterations required for a dominating group to be established is 21. After a further 20 iterations on average, all agents arrive at a consensus heading direction with an average error within 0.05 rad of the informed agents' goal.

Informed agent group size and confidence: Here we investigate the probability that $\mathcal{L}_1$ becomes a dominating group, as a function of initial group size and confidence level. Fig. 3(b) shows the results of varying the number of agents in $\mathcal{L}_1$ from 1 to 9 ($|\mathcal{L}_2|$ correspondingly varies as $9 \rightarrow 1$), and of varying the confidence of the agents in $\mathcal{L}_1$ from 1 to 9 (weights of agents in $|\mathcal{L}_2|$ correspondingly change from 9 to 1). Each data point represents 50 randomly initialized trials. Increasing the number or confidence of informed agents in a group generally yields a higher dominating probability.

Online changes in informed agents: During the consensus process, some initially non-informed agents might have the opportunity to make new observations that convert them into informed agents. Conversely, initially informed agents might lose their informed status, e.g., if they lose track of the group goal. We thus consider the case where the size of $\mathcal{L}_1$ changes while the ILR algorithm is running. Initially, we set $|\mathcal{L}_1| = |\mathcal{L}_2| = 14$ and $w_1 = w_2 = 5$. We randomly select a agent to leave or join $\mathcal{L}_1$ every five time steps. The total number of selected agents to join or leave the group varies from 1 to 13. Fig. 3 (c) shows that a group's probability of dominating (averaged from 200 trials) increases as more agents join the group, and decreases as more agents leave.

Goal variability within groups: In some cases, informed agents that are properly considered part of the same group may have slightly different goal states. For instance, noisy observations or perception difference may prevent different agents observing the same target from setting identical goals. Because ILR makes no reference to

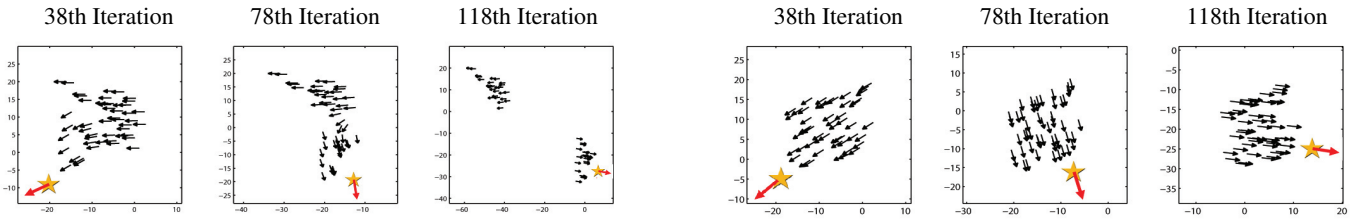

(a) (b)

Figure 4: Snapshots of a multi-agent system (black arrows) performing collective pursuit of a target (star, red arrow). The target changes its heading direction to a new random value every 20 time steps. (a) Agents use Algorithm 1 only. The system fragments into two separated groups after 78 time steps in this run. (b) Agents use Algorithms 1 and 3 (fast decision propagation). The group remains cohesive for 500 time steps, at which point the experiment ended.

Algorithm 3 Pseudocode for agents to achieve fast decision propagation. $\mathcal{N}_i^*$ is the set of neighboring agents whose states differ from x_i by at least $\bar{\phi}$. r is a random number drawn from a uniform distribution between 0 and 1. We choose constants $k_1 = 4 \sim 6$, $k_2 = 2$, $\rho = 0.4$, and $\bar{\phi} = 0.4$. Agents execute this fast decision propagation routine along with the standard algorithm (Alg. 1).

loop
  if $r \leq \frac{|\mathcal{N}_i^*|^{k_1}}{|\mathcal{N}_i^*|^{k_1}+|\mathcal{N}_i - \mathcal{N}_i^*|^{k_2}}$
  and $\max_{\{j,k\}\in\mathcal{N}_i^*} ||x_j - x_k|| < \rho$ **then**
    $x_i^*(t+1) = \frac{1}{|\mathcal{N}_i^*|}\sum_{j|a_j\in\mathcal{N}_i^*} x_j^*(t)$
    if a_i is non-informed **then**
      a_i becomes informed
      $w_i(t+1) = \frac{1}{|\mathcal{N}_i^*|}\sum_{j|a_j\in\mathcal{N}_i^*} w_j(t)$

group identity, such variability need not interfere with the algorithm. Experiments that add zero-mean Gaussian noise to each informed agent's goal state demonstrate that a group's probability of dominating is not significantly affected by moderate variability of this sort (Fig. 3 (d)).

The condition presented in line 3 of Algorithm 2, which evaluates whether a large enough fraction of an agent's neighbors have similar states, is experimentally effective for cases where informed agents are localized into spatial clusters. Other conditions may be more appropriate for other scenarios. For instance, our preliminary results suggest that if informed agents are randomly distributed in space, then the condition $\|x_i(t) - x_i^*\| \leq \delta$—i.e., informed agents increase their confidence if they are able to travel sufficiently close to their goal direction—results in more effective performance.

6. FAST DECISION PROPAGATION

The control strategy we have presented so far requires agents to iteratively approach a group consensus decision. If the iteration time is long, consensus can be slow in coming, since tens of updates may be required. In nature, groups of animals can react very quickly to sudden environmental changes. For example, when a flock of starlings maneuvers to evade a predator's pursuit, starlings on the far side of the flock from the predator react swiftly even without direct perception of the threat. Inspired by a process by which a school of fish arrives at a quorum decision quickly [14], here we further extend our framework in a way that lets a group of agents reach consensus quickly for coping with rapid environmental changes.

The basic concept that enables fast propagation is that we allow non-informed agents to acquire informed status and preferred goal states through interactions with neighbors. In this way the size of

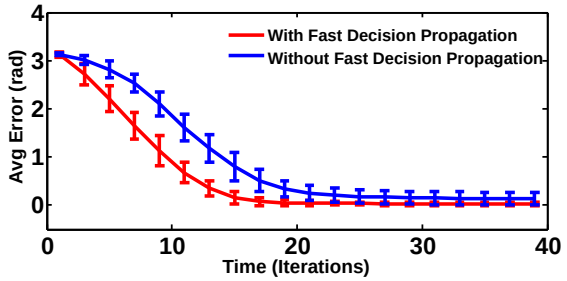

Figure 5: In a task in which all agents are initially aligned with a target that then instantaneously reverses direction, agents consistently reach consensus on the new direction more quickly when using fast decision propagation (Algorithms 1+3, red curve) than when using the implicit leadership algorithm (Algorithm 1, blue curve) alone. Averages are over 100 trials.

the informed group increases. An uninformed agent changes its status to informed when it observes that a significant fraction of its neighbors have goal states different from its own, and that all of those neighbors have goals similar to each other. Already-informed agents can update their goal state x_i^* under the same conditions. Because the convergence speed increases with the number of informed agents (Theorem 4.2), we expect this extension to result in faster consensus.

Algorithm 3 shows the fast decision propagation algorithm. We define $\mathcal{N}_i^*$ as the set of neighbors of a_i whose states differ from x_i by at least some amount $\bar{\phi}$. If all agents in $\mathcal{N}_i^*$ have similar states ($\max_{\{j,k\}\in\mathcal{N}_i^*} ||x_j - x_k|| < \rho$), x_i changes to be the average of the $\mathcal{N}_i^*$ agents' goal states with probability $P = \frac{|\mathcal{N}_i^*|^{k_1}}{|\mathcal{N}_i^*|^{k_1}+|\mathcal{N}_i - \mathcal{N}_i^*|^{k_2}}$. If a_i was non-informed, it also becomes informed, setting its weight to be the average of the $\mathcal{N}_i^*$ agents' weights. The choice of expression for P is based on a study of how fish respond to their neighbors in quorum decision-making for travel direction [14]. In our implementation, we set $k_1 = 4 \sim 6$ and $k_2 = 2$, which yields $P \sim 1$ when the fraction of $\mathcal{N}_i^*$ among a_i's neighbors is greater than 0.5. This formulation also prevents the propagation of a few agents' bad decisions: $P \sim 0$ when $\frac{|\mathcal{N}_i^*|}{|\mathcal{N}_i|} < 0.3$.

Experimental Results: We evaluate Algorithm 3 in a challenging tracking task using the same simulation framework as the other experiments reported above. A target of collective pursuit (Fig. 4) moves distance 0.05 per time step and rotates its heading direction counterclockwise by a random amount between 0.3–0.9 radians every 20 steps. Agents within a distance of 5 can perceive the target, and accordingly act as informed agents. To effectively follow the target's rapidly changing trajectory, the whole group needs to propagate the information of the informed agents as fast as possible.

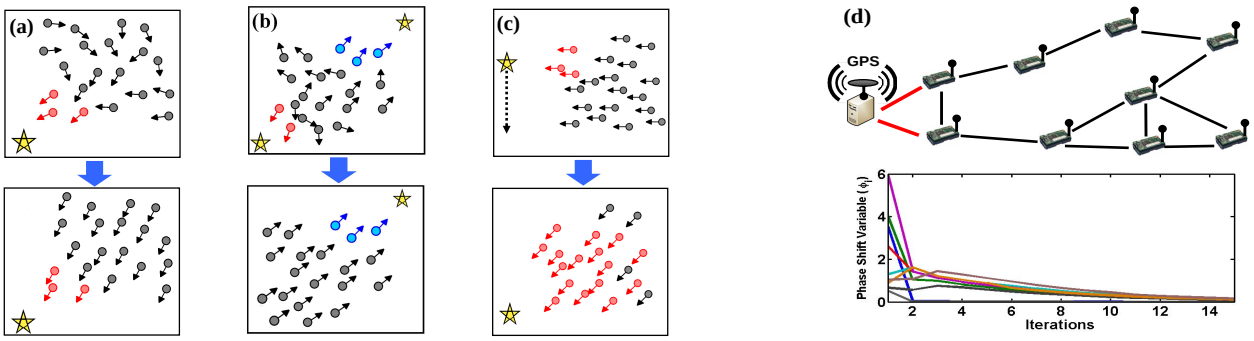

Figure 6: Multi-robot exploration (a–c) and sensor network time synchronization (d) tasks. (a) Red agents localize the target and the whole group iteratively achieves a collective decision. (b) Two groups simultaneously localize two different targets. After running ILR, the blue informed agents dominate the red informed agents and guide the whole group towards their target. (c) The target swiftly changes its direction. Fast decision propagation lets more agents become informed agents to allow the group to track the target quickly. (d) Two sensor nodes that can access UCT are informed agents (top). The whole group eventually achieves the same phase shifts as the informed agents ($\phi_i(t) \rightarrow \phi^* = 0\,\forall i$).

Experiments used 100 different target trajectories of 500 time steps, tested with each of the following two control laws. Agents programmed with the basic implicit leadership algorithm (Algorithm 1) alone fragmented into disconnected groups in *every* trial (75% fragmented in the first 100 steps), while groups of agents using fast decision propagation (Algorithms 1+3) remained connected until the end of the run in 31% of trials. Fig. 4 shows snapshots from one example trial. In (a), we see that a large number of agents end up traveling in a different direction from that of the target, resulting in two spatially separated groups that cannot communicate with each other. By contrast, in (b), the fast decision propagation algorithm much more consistently lets the group reach consensus quickly enough to track the target's movement as a single cohesive unit.

We also test how fast the group can react to a sudden reversal of target direction. Fig. 5 shows the results of experiments from 100 trials in which all agents and the target are initialized with identical direction and random position, and the target's direction is instantaneously reversed by π. Systems using fast decision propagation consistently adjust to the new direction significantly faster than systems using the basic implicit leadership algorithm alone.

7. APPLICATIONS

In this section, we demonstrate applying our framework to three diverse applications. (1) In a multi-robot exploration task, robots search for a target in a 2D space. Robots close enough to the target to localize it act as informed agents to allow the whole group to reach consensus on moving toward the target. (2) In a sensor network time synchronization task, sensor nodes close to a base station can access a global clock and thus better serve as informed agents. The whole group can then achieve more accurate clock settings than if a standard information-symmetric distributed consensus approach is used. (3) In a modular terrain-adaptive bridge task, robotic modules at the ends of a chain act as informed agents with their positions determined by the environment, and when the modules in between reach consensus, a straight bridge is formed.

7.1 Multi-Robot Exploration Task

Here we consider a team of mobile robots searching for a target. Each robot agent follows a movement direction $d_i(t)$ based on Reynolds's flocking rule [11]:

$$d_i(t) = \alpha_x \cdot x_i(t) + \alpha_c \cdot d_c(t) + \alpha_s \cdot d_s(t)$$

where x_i is the heading direction computed by our control rule Algorithm 1; $\alpha_x, \alpha_c, \alpha_s$ are constants; d_c and d_s are functions of the positions of neighboring agents, with d_c associated with maintaining a minimum distance to avoid collisions and d_s associated with attraction between more distant agents. At each time step, each agent communicates $x_i(t)$ to its neighbors.

Fig. 6 (a) shows how agents explore the area with different heading directions (indicated as arrows) until several agents sense the target. The agents set their goal states as the direction pointing towards the target, and their confidence w_i to positive values. All agents proceed to reach a group decision of moving toward the target. Fig. 6 (b) shows how ILR allows all agents to reach a single consensus when two different targets are located simultaneously. If the target is mobile, the fast decision propagation algorithm allows the group to track the target more effectively (Fig. 6 (c)).

7.2 Sensor Network Time Synchronization

In wireless sensor networks, maintaining synchronized clocks between nodes is important. To save power, sensor nodes are usually programmed to stay in sleep mode most of the time. Thus it is important for neighboring nodes to synchronize their timers so they can wake up at the same time and communicate with each other. One strategy is to use a decentralized synchronization protocol inspired by coupled oscillator models in biology [6]. In this model, each node maintains a phase $\theta_i(t) = \omega t + \phi_i(t)$, where $\theta_i(t), \phi_i(t) \in [0, 2\pi]$ ($0 \equiv 2\pi$) and ω is a constant related to frequency of oscillation. $\phi_i(t)$ is the phase shift variable, and each node's phase cycles periodically. When $\theta_i(t) = 2\pi$, the node wakes up and "fires". If all nodes can achieve identical $\phi_i(t)$, the network is synchronized. The control law in [6] is similar to that of distributed consensus, and each node constantly exchanges its phase shift variable with its neighbors. The discrete time step version of the protocol can be written as:

$$\phi_i(t+1) \leftarrow \phi_i(t) + \alpha \sum_{a_j \in \mathcal{N}_i} (\phi_j(t) - \phi_i(t))$$

where α is a small constant.

One major drawback of this approach is that it fails to take into account that some sensor nodes' timers are more accurate than others. In real-world sensor network deployment, some nodes close to a base station can access and synchronize to Coordinated Universal Time (UCT) via GPS. With our framework, we set those nodes' desired phase shift to match the UCT timer, denoted as ϕ^*, and their confidence to be positive. The control law then becomes:

$$\phi_i(t+1) \leftarrow \frac{1}{1+w_i} \cdot \phi_{\text{avg}}(t) + \frac{w_i}{1+w_i} \cdot \phi^*$$

where $\phi_{\text{avg}}(t)$ is the average phase shift of agent i and its neighbors.

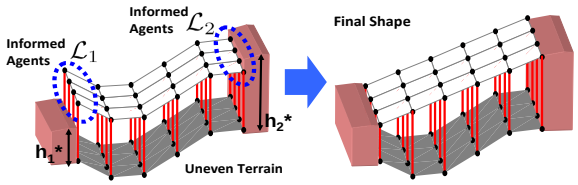

Figure 7: Modular robot shape formation task. Modules at the two ends form two informed agent groups ($\mathcal{L}_1$ and $\mathcal{L}_2$). Agents in $\mathcal{L}_1$ and $\mathcal{L}_2$ set their goal states to be the heights of the adjacent cliffs (h_1^* and h_2^*). The robot converges to a shape forming a straight slope (Right).

Fig. 6 (d) shows an example scenario for this time synchronization task. The two leftmost nodes can access UCT and thus act as informed agents. The bottom figure shows the phase shift variables to which each node converges, which eventually match those of the informed agents ($\phi_i(t) = \phi^* = 0\,\forall i$).

7.3 Modular Robot Shape Formation Task

In [15], the authors describe the design and control of modular robots that adapt to the environment. One example is a modular bridge, which is composed of actuated modular "legs" connected to each other by a flexible bridge surface. When the bridge is placed on an unknown terrain, each modular leg can sense the tilt angles of the local surface elements and adaptively change its height until a level surface is achieved. They show that a distributed consensus strategy can be used to solve this problem. One problem with the approach presented in [15] is that while it allows all the modules to arrive at the same height, it cannot be applied to cases where a subset of agents need to maintain specific desired heights.

Using the approach presented here, one can easily interact with this system to incorporate desired heights; e.g., by simply fixing the height of one of the modules, the entire system will converge to agreement at that height if feasible. The control law can be written:

$$h_i(t+1) \leftarrow \frac{1}{1+w_i}\left(h_i(t) + \alpha \sum_{a_j \in \mathcal{N}_i} \theta_{ij}\right) + \frac{w_i}{1+w_i} h_i^*$$

where $h_i(t)$ and h_i^* represent current and desired heights of the agents (modules) and θ_{ij} represents the tilt sensor reading between agents i and j. By fixing the end agents at two different heights we can further create smooth inclines. Fig. 7 shows a simulation of the modular bridge where the agents at the two ends are fixed to the respective heights of the cliffs; the system converges to a smooth incline (Theorem 4.3). This type of implicit leadership provides a simple way for influencing and interacting with the modular robotic system and equally applies to the other applications presented in [16] such as the pressure-adaptive column and gripper.

8. CONCLUSIONS

In this work we have presented an implicit leadership algorithm based on a simple, purely local control law. This approach allows all agents to follow a single rule and efficiently reach a common group decision without complex coordination mechanisms. We have characterized the algorithm's theoretical properties both analytically and experimentally, and presented extensions that handle multiple conflicting goals among informed agents and achieve fast decision propagation. In a distributed multi-agent system, it can be tedious to establish infrastructure and coordination mechanisms for agents to integrate unpredictably localized information to make collective decisions. The biologically inspired framework presented here provides a simple alternative that allows the group to reach consensus regarding distributed and dynamic information.

The control rule presented in Algorithm 1 can be extended in various ways; implicit leadership reinforcement and fast decision propagation (Algorithms 2 and 3) can be seen as two such extensions. Future work may explore further variations, which additionally may be used in combination; e.g., while we considered Algorithms 2 and 3 separately above, both can be applied to the same multi-agent system simultaneously. Our preliminary experiments show that systems using both extensions can achieve effective performance in both resolution of conflicting goals and rapid convergence to consensus.

We have shown how our framework can be applied to several diverse application areas. Other applications may likewise benefit from incorporating our approach. Many existing control rules can be modified to incorporate a decision state that is observed or explicitly communicated between agents, as demonstrated with the example of Reynolds's rules for flocking above. Our algorithm can also be seen as a useful tool that can be exploited in various multi-agent coordination tasks, with the advantage that controllers for those tasks thereby acquire the associated convergence guarantees.

Acknowledgement

This research is sponsored by NSF (#0829745) and Wyss Institute.

9. REFERENCES

[1] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, 1999.

[2] I. Couzin, J. Krause, N. Franks, and S. Levin. Effective leadership and decision making in animal groups on the move. *Nature*, 433, 2005.

[3] A. Jadbabaie, J. Lin, and A. Morse. Coordination of groups of autonomous agents using nearest neighbor rules. *IEEE Trans. on Auto. Control*, 2002.

[4] M. Ji. *Graph-Based Control of Networked Systems*. PhD thesis, Georgia Tech, 2007.

[5] N. E. Leonard and E. Fiorelli. Virtual leaders, artificial potentials and coordinated control of groups. In *Proc. IEEE CDC*, 2001.

[6] D. Lucarelli and I. Wang. Decentralized synchronization protocols with nearest neighbor communication. In *Proc. Sensys*, 2004.

[7] M. Dorigo et al. Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, 2004.

[8] R. Olfati-Saber, J. Fax, and R. Murray. Flocking for multi-agent dynamic systems: Algorithms and theory. In *IEEE Trans. on Auto. Control*, 2006.

[9] R. Olfati-Saber, J. Fax, and R. Murray. Consensus and cooperation in networked multi-agent systems. *Proc. IEEE*, Jan 2007.

[10] R. Poli, J. Kennedy, and T. Blackwell. Particle swarm optimization. *Swarm Intelligence*, June 2007.

[11] C. W. Reynolds. Flocks, herds, and schools: a distributed behavioral model. In *Proc. SIGGRAPH*, 1987.

[12] D. Rus, Z. Butler, K. Kotay, and M. Vona. Self-reconfiguring robots. *Comm. of the ACM*, 2002.

[13] H. Tanner and V. Kumar. Leader-to-formation stability. *IEEE Trans. on Robotics and Automation*, 2004.

[14] A. Ward, D. Sumpter, I. Couzin, P. Hart, and J. Krause. Quorum decision-making facilitates information transfer in fish shoals. *PNAS*, May 2008.

[15] C. Yu and R. Nagpal. Sensing-based shape formation tasks on modular multi-robot systems: A theoretical study. In *Proc. AAMAS*, 2008.

[16] C. Yu and R. Nagpal. Self-adapting modular robotics: A generalized distributed consensus framework. In *Proc. ICRA*, 2009.

APPENDIX

A. Proof of Theorem 4.1

We first aggregate all agents' state update dynamics into a single collective dynamic system. Here we consider the case of a single informed agent group $\mathcal{L}_1$ with common goal state x^*. Add an auxiliary agent a^* assumed to connect to all informed agents. Agent states can be aggregated into an $(n+1)$-dimensional vector:

$$X(t) = [x_1(t), \cdots, x_n(t), x^*]^T$$

The system's collective dynamics can be written as:

$$X(t+1) = A(t) \cdot X(t) = \prod_{m=1}^{t} A(m)X(1) \quad (1)$$

$$= \prod_{m=1}^{t} \begin{pmatrix} F(m) & H(m) \\ 0 & 1 \end{pmatrix} X(1) = \begin{pmatrix} P(t) & Q(t) \\ 0 & 1 \end{pmatrix} X(1)$$

where $F(m)$ is a $n \times n$ matrix, with $F_{ij}(m) = \frac{1}{(1+w_i)|\mathcal{N}_i(m)+1|}$ if $a_j \in \mathcal{N}_i(m)$ and $a_j \in \mathcal{L}_1$, $F_{ij}(m) = \frac{1}{|\mathcal{N}_i(m)+1|}$ if $a_j \in \mathcal{N}_i(m)$ and $a_j \notin \mathcal{L}_1$, and $F_{ij}(m) = 0$ otherwise. $\mathcal{N}_i(m)$ indicates the set of agent i's neighbors at time m. $H(m)$ is a n-dimensional column vector, $H_i(m) = \frac{w_i}{(1+w_i)}$ if $a_i \in \mathcal{L}_1$ and $H_i(m) = 0$ otherwise. We further define:

1. Matrix maximum norm: $\|M\|_\infty = \max_i \sum_j M_{ij}$;
2. Matrix minimum norm: $\|M\|_{-\infty} = \min_i \sum_j M_{ij}$;
3. $d_{\max} = \max_i \text{dist}(a_i, n_i^*) \forall a_i \notin \mathcal{L}_1$, the maximal distance between a non-informed agent a_i and its closest informed agent n_i^*.

Static Topology: In this case, $A(m)$, $F(m)$, and $H(m)$ are time-invariant. We can let $A(m) = \bar{A}$, $F(m) = \bar{F}$, $H(m) = \bar{H}$, and

$$X(t+1) = \prod_{m=1}^{t} \bar{A}X(1) = \prod_{m=1}^{t} \begin{pmatrix} \bar{F} & \bar{H} \\ 0 & 1 \end{pmatrix} X(1) \quad (2)$$

Since $\bar{A}$ is a row stochastic matrix all of whose main-diagonal elements are positive, all elements in $\bar{A}$ are nonnegative ($\bar{A}_{ij} \geq 0$) and $\sum_j \bar{A}_{ij} = 1 \, \forall i$. We can show that $(\bar{F}^{d+1})_{i,j} > 0$ for any $\{i, j\}$ that are d hops apart, and the connectivity graph G is connected. Since $Q(t) = \sum_{m=1}^{t} \bar{F}^{m-1}\bar{H}$ and $\bar{H}_i > 0 \forall i | a_i \in \mathcal{L}_1$, we can ensure $Q_i(t) > 0 \, \forall i$ after $d_{\max} + 2$ time steps. Therefore, $\|Q(t)\|_{-\infty} > 0$. Since $\|P(t)\|_\infty = 1 - \|Q(t)\|_{-\infty}$, we can derive $0 < \|P(t)\|_\infty < 1$ after $d_{\max} + 2$ time steps. The maximum norm of $P(t)$ decreases at least every $d_{\max} + 2$ time steps. Thus, $\lim_{t\to\infty} \|P(t)\|_\infty = 0$ and $\lim_{t\to\infty} P(t) = \mathbf{0}, \lim_{t\to\infty} Q(t) = \mathbf{1} \Rightarrow \lim_{t\to\infty} x_i(t) = x^* \, \forall i$.

Dynamic Topology: We define $\bar{G} = \bigcup_{t'}^{t'+\Delta} G(t)$ as a union of connectivity graphs over a finite time period Δ. If agents' topology is informed agent-connected during a finite period τ, then for each non-informed agent a_i, there exist at least $d_{\max}+2$ instances within τ of a $\bar{G}$ that connects a_i with some informed agent. Thus, there are at least $d_{max} + 2$ instances of $A' = \prod_{t'}^{t'+\Delta} A(t)$ within τ.

We expand $Q(\tau) = \sum_{m=1}^{\tau} \left(\prod_{k=1}^{m-1} F(k)\right) H(m)$. Since all elements on the main diagonal of $A(t)$ are always positive and there are at least $d_{max} + 2$ instances of A', we can derive $Q_i(t) > 0 \, \forall i$ and $t > \tau$. Following a similar procedure as in the static topology case, we can show $0 < \|P(t)\|_\infty < 1$ after τ time steps and the maximum norm of $P(t)$ contracts. Thus,

$$\lim_{t\to\infty} P(t) = \mathbf{0}, \lim_{t\to\infty} Q(t) = \mathbf{1} \Rightarrow \lim_{t\to\infty} x_i(t) = x^* \, \forall i$$

B. Proof of Theorem 4.2

We first demonstrate that the system's convergence speed increases when a previously non-informed agent becomes informed, increasing the number of informed agents from m to $m+1$. We first show this is true in the static topology case. As shown in the proof of Theorem 4.1, the speed at which the maximum norm $\|P(t)\|_\infty$ approaches zero determines the system convergence speed. We start by adding one additional informed agent to the system. Let $\mathcal{L}_1' = l \cup \mathcal{L}_1$ denote the new informed agent group after adding agent l to the original $\mathcal{L}_1$. We use $\bar{F}'$, $P'(t)$, $Q'(t)$ to denote new agent matrices for the system with informed agent group $\mathcal{L}_1'$.

Let $Q_l(t)$ and $Q_l'(t)$ be elements of Q corresponding to l, in the m-informed agent and $(m+1)$-informed agent cases respectively. We can derive matrix elements at $t = 2$ from Eq. 2:

$$Q_l'(2) = \left(\sum_j \bar{F}_{lj}' Q_j'(1)\right) + \frac{w_l}{1+w_l} > \sum_j F_{lj}Q_j(1) = Q_l(2)$$

With the same procedure, we can show $Q_k'(3) > Q_k(3)$ for all agents k that are neighbors of l, $Q_r'(4) > Q_r(4)$ for all r that are neighbors of k, etc. Thus $Q_i'(d'+1) > Q_i(d'+1) \, \forall i$, where $d' = \max_i \text{dist}(l, i)$.

Thus, $\|P'(d'+1)\|_\infty < \|P(d'+1)\|_\infty$ and agents converge faster in the $\mathcal{L}_1'$ case. Adding multiple informed agents at the same time is a trivial generalization of this proof.

A similar proof applies to the case of changing weight(s) of informed agent(s) from w_l to $w_l' > w_l$. Let $Q_l'(t)$ be an element of Q after the weight increase, and we can write $Q_l'(2) = \left(\sum_j \bar{F}_{lj}' Q_j'(1)\right) + \frac{w_l'}{1+w_l'} > \left(\sum_j F_{lj}Q_j(1)\right) + \frac{w_l}{1+w_l} = Q_l(2)$. We can then show $Q_k'(3) > Q_k(3) \, \forall k \in \mathcal{N}_l$, etc. After $d'+1$ time steps, we have $Q_i'(d'+1) > Q_i(d'+1) \forall i$.

With dynamic topology, if there exists a path between a_l and *all* non-informed agents, we can similarly show that $Q_l'(2) = \left(\sum_j \bar{F}_{lj}' Q_j'(1)\right) + \frac{w_l}{1+w_l} > \sum_j F_{lj}Q_j(1) = Q_l(2)$. If such a path exists $d'+1$ times in a finite period τ, we can show $Q_i'(\tau) > Q_i(\tau) \forall i$ and $\|P'(\tau)\|_\infty < \|P(\tau)\|_\infty$. Therefore, adding one more informed agent makes the system converge faster.

C. Proof of Theorem 4.3

With q different informed agent groups, we introduce auxiliary agent states $\bar{x}_1^* \cdots \bar{x}_q^*$, with agents in each informed agent group $\mathcal{L}_i$ connected to auxiliary agent with state $\bar{x}_i^*$. We can formulate the collective dynamics of all agents as Eq. 2 of Theorem 4.1 with a different generating matrix:

$$\bar{A} = \begin{pmatrix} \bar{F} & \bar{H} \\ 0 & I \end{pmatrix}, \; \prod_{m=1}^{t} \bar{A} = \bar{A}^t = \begin{pmatrix} P(t) & Q(t) \\ 0 & I \end{pmatrix}$$

where $\bar{F}$ and $P(t)$ are $n \times n$ matrices, and $\bar{H}$ and $Q(t)$ are $n \times q$ matrices. Following the same procedure as for Thm. 4.1, we can show that the maximum matrix norm of $P(t)$ approaches zero:

$$\lim_{t\to\infty} \|P(t)\|_\infty = 0 \text{ and } \sum_j Q_{ij}(t) = 1 \forall i$$

Since $Q(t)$ is multi-column, we show that each element in $Q(t)$ converges to a stable state: $Q_{ij}(t) \to Q_{ij}^* \, \forall i, j$. We decompose

$$\bar{A}^t = \begin{pmatrix} P(t) & Q(t) \\ 0 & I \end{pmatrix} = \begin{pmatrix} P(t-1) & Q(t-1) \\ 0 & I \end{pmatrix} \cdot \begin{pmatrix} \bar{F} & \bar{H} \\ 0 & I \end{pmatrix}$$

Since the bottom right submatrix of $\bar{A}$ is a $q \times q$ identity matrix, we can get $Q_{ij}(t) \geq Q_{ij}(t-1) \, \forall i, j$; and $Q_{ij}(t)$ is monotonically nondecreasing with t for all i, j. Therefore, as $\sum_j Q_{ij}(t)$ approaches $\mathbf{1}$, $Q_{ij}(t)$ will also approach a fixed value $Q_{ij}^* \, \forall i, j$. Since each informed agent group's state is different, each agent converges to a potentially different fixed state: $\lim_{t\to\infty} x_i(t) = \sum_j Q_{ij}^* \bar{x}_j^*$.