Laplacian Sheep: A Hybrid, Stop-Go Policy for Leader-Based Containment Control

G. Ferrari-Trecate[1,2], M. Egerstedt[3], A. Buffa[4], and M. Ji[3]

[1] INRIA, Domaine de Voluceau, Rocquencourt - B.P.105,
78153, Le Chesnay Cedex, France
[2] Dipartimento di Informatica e Sistemistica,
Universitá degli Studi di Pavia, Via Ferrata 1, 27100 Pavia, Italy
giancarlo.ferrari@unipv.it
[3] Georgia Institute of Technology,
School of Electrical and Computer Engineering,
Atlanta, GA 30332, USA
{magnus, mengji}@ece.gatech.edu
[4] Istituto di Matematica Applicata e Tecnologie Informatiche, C.N.R.,
Via Ferrata, 1, 27100 Pavia, Italy
annalisa@imati.cnr.it

Abstract. The problem of driving a collection of mobile robots to a given target location is studied in the context of partial difference equations. In particular, we are interested in achieving this transfer while ensuring that the agents stay in the convex polytope spanned by dedicated leader-agents, whose dynamics will be given by a hybrid Stop-Go policy. The resulting system ensures containment through the enabling result that under a Laplacian, decentralized control strategy for the followers, these followers will converge to a location in the convex leader polytope, as long as the leaders are stationary and the interaction graph is connected. Simulation results testify to the viability of the proposed, hybrid control strategy.

1 Introduction

This paper investigates a particular subarea of multi-agent control, namely the so-called *containment* problem. (See for example [1, 2].) The problem is to drive a collection of autonomous, mobile agents to a given target location while guaranteeing that their motion satisfies certain geometric constraints. These constraints are there to ensure that the agents are contained in a particular area during their transportation. These types of issues arise for example when a collection of autonomous robots are to secure and then remove hazardous materials. This removal must be secure in the sense that the robots should not venture into populated areas or in other ways contaminate their surroundings.

We approach this problem from a leader-follower point-of-view [3, 4, 5]. In particular, we will let the agents move autonomously based on local, consensus-like interaction rules, commonly found in the literature under the banner of

João Hespanha and A. Tiwari (Eds.): HSCC 2006, LNCS 3927, pp. 212–226, 2006.

algebraic graph theory [6, 7, 8]. The reason for this is that the robot motion can be described using the graph Laplacian. However, we will augment this control structure with the addition of leader-agents. These leaders are to define vertexes in a convex polytope (the leader-polytope) and they are to move in such a way that the target area is reached while ensuring that the follower-agents stay in the convex polytope spanned by the leaders, as shown in Fig. 1. As such, the followers movements are calculated in a decentralized manner according to a fixed interaction topology, while the leaders are assumed to be able to detect if any of the followers violate the containment property. This strategy also explains the title "Laplacian Sheep" since the followers are moving using a "Laplacian-based" control strategy. However, they are to be "herded" like sheep by the leaders, and hence the title.

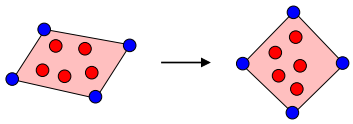

Fig. 1. The containment problem: The leaders are to move in such a way that the followers remain in the convex leader-polytope for all times

It should be noted already at this point that although the subject matter is multi-agent control, our proposed solution to the problem of selecting the leader dynamics will be hybrid. In particular, we will use a Stop-Go policy [9, 10], in which the leaders move according to a decentralized formation control strategy until the containment property is about to be violated. At this point, they stop and let the followers settle back into the leader-polytope before they start moving again. For such a strategy to be successful, a number of results are needed, including a guarantee that the Laplacian-based follower-control will in fact drive the followers back into the leader-polytope. Moreover, we must also ensure that such a control strategy is feasible in the sense of non-Zeno, live in the sense of not staying in the Stop mode indefinitely, and convergent in the sense that the target area is in fact reached. These are the main issues under investigation in this paper.

In order to properly understand the behavior of such a control system, some initial results in multi-agent control are needed. We will use the recently developed framework of Partial difference Equations (PdEs) [11] for this, and, in particular, we will show that as long as the interaction graph is connected and the leaders are stationary, the followers will always converge to locations in the convex hull spanned by the leaders. This result enables the use of a Stop-Go policy since by halting the evolution of the leaders, the containment property can be ensured.

The outline of this paper is as follows: In Section 2 we present the mathematical preliminaries needed to prove the main convergence result for the case with stationary leaders, in Section 3. The hybrid Stop-Go control policy is given in Section 4, followed by an illustrative example, in Section 5. Additional

extensions to the proposed control strategy, including a hierarchical layering of the formation, are given in Section 6.

2 Background and Mathematical Preliminaries

Even though the main focus of this paper is the development of hybrid control strategies for the leader-agents in charge of "herding" the followers, this task will rely on a collection of enabling results in (non-hybrid) multi-agent control. These enabling results will allow us to structure the hybrid Stop-Go controller in such a way that containment is achieved. Hence, before we can start defining any hybrid control laws, some room must be given to multi-agent control. In this section we will present the basic mathematical framework, and the main containment results will follow in the next section.

We start summarizing basic notions of graph theory. For more details we defer the reader to [12]. An undirected graph G is defined by a set $\mathcal{N}_G = \{1, \dots N\}$ of *nodes* and a set $\mathcal{E}_G \subset \mathcal{N}_G \times \mathcal{N}_G$ of *edges*. We will also use $|\mathcal{N}_G|$ for denoting the cardinality of $\mathcal{N}_G$. Two nodes x and y are *neighbors* if $(x, y) \in \mathcal{E}_G$. The neighboring relation is indicated with $x \sim y$ and $\mathcal{P}(x) = \{y \in \mathcal{N}_G : y \sim x\}$ collects all neighbors to the node x. A *path* $x_0 x_1 \dots x_L$ is a finite sequence of nodes such that $x_{i-1} \sim x_i$, $i = 1, \dots, L$. A graph G is *connected* if there is a path connecting every pair of distinct nodes. G is *complete* if $\mathcal{E}_G = \mathcal{N}_G \times \mathcal{N}_G$.

Definition 1. *Let $S = (\mathcal{N}_S, \mathcal{E}_S)$ be an undirected host graph and $\mathcal{N}_{S'} \subset \mathcal{N}_S$. The subgraph S' associated with $\mathcal{N}_{S'}$ is the pair $(\mathcal{N}_{S'}, \mathcal{E}_{S'})$ where $\mathcal{E}_{S'} = \{(x, y) \in \mathcal{E}_S : x \in \mathcal{N}_{S'}, y \in \mathcal{N}_{S'}\}$.*

Definition 1 allows basic operations in set theory to be extended to graphs.

Definition 2. *Let S_1 and S_2 be to subgraphs of the graph S. Then, $S_1 \cup S_2$, $S_1 \cap S_2$, $S_1 \backslash S_2$ are the graph associated with $\mathcal{N}_{S_1} \cup \mathcal{N}_{S_2}$, $\mathcal{N}_{S_1} \cap \mathcal{N}_{S_2}$, and $\mathcal{N}_{S_1} \backslash \mathcal{N}_{S_2}$, respectively.*

For our purposes, we will often use graphs with a boundary.

Definition 3. *Let S be a subgraph of G. The boundary of S is the subgraph $\partial S \subset G$ associated with $\mathcal{N}_{\partial S} \doteq \{y \in \mathcal{N}_G \setminus \mathcal{N}_S \ : \ \exists x \in \mathcal{N}_S \ : \ x \sim y\}$. The closure of S is $\bar{S} = \partial S \cup S$.*

Note that the definition of the boundary of a graph depends upon the host graph G. This implies that if one considers three graphs $S' \subset S \subset G$, the boundaries of S' in S and in G may differ.

In our case, the nodes of the host graph G represent agents and the edges are communication links. In particular, an agent x has access to the states of all its neighbors and can use this piece of information to compute its control law. In this setting, partial communication amounts to considering incomplete graphs. However, we always assume that the host graph is connected, otherwise the agents are split in one or more sub-groups that do not exchange information.

In order to model the collective behavior of the agents we will use functions $f : \mathcal{N}_G \mapsto \mathbb{R}^d$ defined over a graph G [13]. The *partial derivative* of f is defined as $\partial_y f(x) \doteq f(y) - f(x)$ and enjoys the following properties: (i) $\partial_y f(x) = -\partial_x f(y)$, (ii) $\partial_x f(x) = 0$ and (iii) $\partial_y^2 f(x) = -\partial_y f(x)$. The Laplacian of f is given by

$$\Delta f(x) \doteq -\sum_{y \in \mathcal{N}_G, y \sim x} \partial_y^2 f(x) = +\sum_{y \in \mathcal{N}_G, y \sim x} \partial_y f(x), \tag{1}$$

where the last identity follows from property (iii). The integral and the average of f are defined, respectively, as

$$\int_G f \, dx \doteq \sum_{x \in \mathcal{N}_G} f(x), \quad \langle f \rangle \doteq \frac{1}{|\mathcal{N}_G|} \int_G f \, dx. \tag{2}$$

Let $L^2(G|\mathbb{R}^d)$ be the Hilbert space composed by all functions $f : \mathcal{N}_G \mapsto \mathbb{R}^d$ endowed with the norm $\|f\|_{L^2}^2 = \int_G \|f\|^2$. We will use the shorthand notation L^2 when there is no ambiguity on the underlying domain and range of the functions.

Let S be a subgraph of G and ∂S its boundary in G, such that $S \cup \partial S = G$ As in [13], we also consider the Hilbert space $H_0^1(S) = \{f \in L^2(G) \ : \ f_{|\partial S} = 0\}$ (see [13] for the definition of a suitable norm on $H_0^1(S)$). Note that a function $f \in H_0^1(S)$ is defined on $\bar{S}$ and possibly non null only on S.

The next theorem, proved in [13], characterize the eigenstructure of the Laplacian operator defined on $H_0^1(S)$.

Theorem 1. *Let G be a connected graph and S a proper subgraph of G. Then, the operator $\Delta : H_0^1(S|\mathbb{R}^d) \mapsto L^2(\bar{S}|\mathbb{R}^d)$ has $|\mathcal{N}_S|d$ strictly negative eigenvalues. Moreover, the corresponding eigenfunctions form a basis for $H_0^1(S|\mathbb{R}^d)$.*

3 Multiple Stationary Leaders

In this section we use PdEs for modeling and analyzing a group of agents with multiple leaders. A leader is just a vehicle that moves along a prescribed trajectory, independently of the motion of all other vehicles. However, followers that are neighbors to the leader can use the leader state in order to compute their control inputs.

Let $r(x,t)$ be the position of the agent x at time $t \geq 0$, where[1] $r \in L^2$. The communication network is represented by the undirected and connected graph G. For distinguishing between leaders and followers, we consider two subgraphs S_F and S_L of G such that $S_F \cup S_L = G$ and $\partial S_F = S_L$, where the subscripts denote "Leaders" and "Followers" respectively. Note that we assume that all agents are either designated as leaders or followers.

As already mentioned in the introduction, we will assume that the followers obey the simple dynamics $\dot{r}(x,t) = u(x,t)$, where

$$u(x,t) \doteq \Delta r(x,t) \tag{3}$$

[1] For sake of conciseness, for a function $f(x,t) : \mathcal{N}_G \times \mathbb{R}^+ \to \mathbb{R}^d$ we will often write $f \in L^2$ instead of $f(\cdot,t) \in L^2$.

is the *Laplacian* control law. Let $\hat{r}(x,t)$, $x \in \mathcal{N}_{\partial S_F}$ (i.e. in the set of leaders, $\mathcal{N}_{S_L}$) be the trajectory of the leaders. Then, the collective dynamics is represented by the model

$$\dot{r}(x,t) = \Delta r(x,t) \qquad x \in \mathcal{N}_{S_F} \tag{4a}$$

$$r(x,t) = \hat{r}(x,t) \qquad x \in \mathcal{N}_{\partial S_F} \tag{4b}$$

endowed with the initial conditions $r(\cdot,0) = \tilde{r} \in L^2(S_F)$.

Model (4) is an example of a continuous-time Partial difference Equation (PdE) with non-homogeneous Dirichlet boundary conditions. We defer the reader to [13, 11, 14] for an introduction to PdEs.

Laplacian control has been one of the most studied control paradigms for multi-agent systems. The main reason is that Laplacian control allows the agents to achieve globally coordinated behaviors, despite its decentralized nature. The main results on Laplacian control available in the literature and specialized to model (4) are:

- in the leaderless case (i.e. $\partial S_F = \emptyset$), the Laplacian control solves the rendez-vous problem, i.e. $r(x,t) \to r^* \in \mathbb{R}^d$, $\forall x \in \mathcal{N}_G$ as $t \to +\infty$. Moreover, the agents achieve *average* consensus, i.e. $r^* = \langle \tilde{r} \rangle$. These results have been established in [15, 16] through the joint use of tools in control theory and algebraic graph theory. A formal analysis of the PdE (4a) has been conducted in [11, 14] showing a complete accordance with results available within the theory of the heat equation [17];
- in the case of a single leader (i.e. $\mathcal{N}_{\partial S_F} = \{x_L\}$) with fixed position (i.e. $\hat{r}(x_L,t) = \bar{r} \in \mathbb{R}^d$), Laplacian control solves the rendez-vous problem with $r^* = \bar{r}$ [15]. This property has also been shown in [11, 14] within the PdE framework, thus highlighting the profound links between model (4) and the heat equation with Dirichlet boundary conditions [17].

A first aim of this paper is to characterize the asymptotic behavior of the followers in the presence of multiple leaders with fixed positions. To this end, for the remainder of this section, we will assume that $\hat{r}(x,t) = \bar{r}(x) \in L^2(\partial S_F)$. The equilibria of (4) are then given by the solutions to the PdE

$$\Delta h(x) = 0 \qquad x \in \mathcal{N}_{S_F} \tag{5a}$$

$$h(x) = \bar{r}(x) \quad x \in \mathcal{N}_{\partial S_F} \tag{5b}$$

and they have been studied in [13]. In particular, [13, Theorem 3.5] shows that if the hosting graph G is connected and $\mathcal{N}_{\partial S_F} \neq \emptyset$ then, the PdE (5) has a unique solution[2] $h(x)$. By analogy with the jargon of Partial Differential Equations, h is termed the *harmonic extension of the boundary conditions* $\bar{r}$.

Our next aim is to verify that $r \to h$ as $t \to +\infty$. Let us consider the decomposition

$$r(x,t) = r_0(x,t) + h(x), \quad r_0 \in H_0^1(S_F) \tag{6}$$

[2] [13, Theorem 3.5] assumes that the subgraph S is *induced* (see [13] for the definition of induced subgraphs). However, a careful examination of the proof, reveals that this assumption is unnecessary.

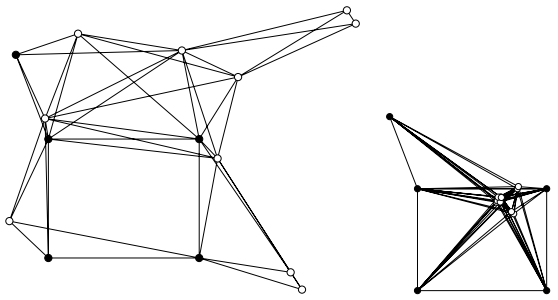

(a) (b)

Fig. 2. An example of the application of Theorem 2 is given. Initially, some of the followers (white) are located outside Ω_L but after a while they have all reached Ω_L, spanned by the stationary leaders (black). The edges between agents capture the information flow in this static interaction graph.

Since h does not depend upon time and $\Delta h = 0$, $\forall x \in \mathcal{N}_{S_F}$, the PdE (4) is equivalent to the following one

$$\dot{r}_0(x,t) = \Delta r_0(x,t) \qquad x \in \mathcal{N}_{S_F} \tag{7a}$$

$$r_0(x,t) = 0 \qquad x \in \mathcal{N}_{\partial S_F} \tag{7b}$$

From (6), it is apparent that the problem of checking if $r \to h$ as $t \to +\infty$ can be recast into the problem of studying the convergence to zero of the solutions to the PdE (7). The fact that $r_0 \to 0$ as $t \to +\infty$ follows from Theorem 1 and it can be shown by proceeding exactly as in the proof of [11, Theorem 7][3].

The next Theorem, proved in Appendix A, highlights a key geometrical feature of $h(x)$. For a set X of points in $\mathbb{R}^d$, $\mathrm{Co}(X)$ will denote its convex hull. Moreover, the set Ω_L is the convex hull of leaders positions, i.e. $\Omega_L \doteq \mathrm{Co}(\{\bar{r}(y),\ y \in \mathcal{N}_{\partial S_F}\})$.

Theorem 2. *Let S_1 be a nonempty connected subgraph of S_F and ∂S_1 be its boundary in G. Then, $\forall x \in \mathcal{N}_{S_1}$ it holds*

$$h(x) \in \mathrm{Co}(\{h(y),\ y \in \mathcal{N}_{\partial S_1}\}). \tag{8}$$

Moreover, one has that $h(x) \in \Omega_L$, i.e. that the position of each follower lies in the convex hull of the leaders positions. Finally, if Ω_L is full dimensional, then $h(x) \in \Omega_L \backslash \partial\Omega_L$, $\forall x \in \mathcal{N}_{S_F}$.

This result is illustrated in Fig. 2.

Another geometrical feature which we need is the following:

Theorem 3. *Suppose that Ω_L is fully dimensional and that $r(x,t)$ is evolving according to (4). Suppose that, at a given time $t = \overline{t}$, there is an agent $x \in \mathcal{N}_{S_F}$ such that $r(x,\overline{t}) \in \partial\Omega_L$. Then, two situations may occur:*

[3] Actually, [11, Theorem 7] proves a stronger property, namely that the origin of (7) is "exponentially stable on the space $H_0^1(S)$". The definition of stability of equilibria on subspaces is provided in [11].

1. *there exists an (affine) hyperplane* χ *such that*

$$r(x,\overline{t}) \in \chi \cap \partial\Omega_L, \text{ and } r(y,\overline{t}) \in \chi \cap \partial\Omega_L \ \forall y \in \mathcal{P}(x).$$

Then:

$$\exists\, \alpha > 0 \ : \ r(x,\overline{t}) + \alpha \dot{r}(x,\overline{t}) \in \chi \cap \partial\Omega_L, \tag{9}$$

2. *otherwise,*

$$\exists\, \alpha > 0 \ : \ r(x,\overline{t}) + \alpha \dot{r}(x,\overline{t}) \in \Omega_L \setminus \partial\Omega_L. \tag{10}$$

Note that (9) means that the velocity of x will be along the hyperplane χ (in other words, the agent may slide on the boundary $\partial\Omega_L$), whereas (10) means that the velocity of x is pointing inside the polytope Ω_L.

Proof: (Theorem 3)
Since $r(x,t)$ obeys to (4), by rearranging terms we obtain:

$$\dot{r}(x,\overline{t}) = -|\mathcal{P}(x)| r(x,\overline{t}) + \sum_{y \in \mathcal{P}(x)} r(y,\overline{t}).$$

Then, setting $\alpha = |\mathcal{P}(x)|^{-1}$, it holds:

$$r(x,\overline{t}) + \alpha \dot{r}(x,\overline{t}) = |\mathcal{P}(x)|^{-1} \sum_{y \in \mathcal{P}(x)} r(y,\overline{t}),$$

i.e., $r(x,\overline{t}) + \alpha \dot{r}(x,\overline{t})$ is the barycenter $b(\mathcal{Y}_x)$ of the polytope $\mathcal{Y}_x \doteq \mathrm{Co}(\{r(y,\overline{t}),\ y \in \mathcal{P}(x)\})$. Note that: first $\mathcal{Y}_x \in \Omega_L$, second, thanks to convexity, the barycenter of $\mathcal{Y}_x$ lies in the relative interior of $\mathcal{Y}_x$. Thus, if all $y \in \mathcal{P}(x)$ verify that $r(y,\overline{t}) \in \chi \cap \partial\Omega_L$ then $\mathcal{Y}_x \subset \chi \cap \partial\Omega_L$ and so does $b(\mathcal{Y}_x)$, i.e. $b(\mathcal{Y}_x) \in \chi \cap \partial\Omega_L$; otherwise $b(\mathcal{Y}_x) \in \Omega_L \setminus \partial\Omega_L$. ■

4 Hybrid Containment Control

Since the motion of the followers is governed directly by the Laplacian control in (4a), what the system designers have control over is the motion of the leaders. In particular, we would like to endow the leaders with a motion that requires as little information sharing as possible while still ensuring containment. For this, we define two distinctly different control modes for the evolution of the leaders. The first of the two control modes is the Stop mode. As the name indicates, this mode corresponds to the leaders halting their movements altogether in order to prohibit a break in the containment:

$$\begin{aligned}
&STOP: \\
&\dot{r}(x,t) = \Delta r(x,t) && x \in \mathcal{N}_{S_F} && \text{(11a)}\\
&r(x,t) = \hat{r}(x,t) && x \in \mathcal{N}_{\partial S_F} && \text{(11b)}\\
&\dot{\hat{r}}(x,t) = 0 && x \in \mathcal{N}_{\partial S_F} && \text{(11c)}
\end{aligned}$$

It is clear that in order to execute this mode, no information is needed for the leaders whatsoever.

The second control mode under consider is the Go mode, in which the leaders move toward a given target location/formation/shape. A number of different control laws can be defined for this, but we, for the sake of conceptual unification, let the Go mode be given by a Laplacian-based control strategy as well.

$$
\begin{aligned}
&GO: \\
&\dot{r}(x,t) = \Delta r(x,t) && x \in \mathcal{N}_{S_F} && \text{(12a)}\\
&r(x,t) = \hat{r}(x,t) && x \in \mathcal{N}_{\partial S_F} && \text{(12b)}\\
&\dot{\hat{r}}(x,t) = \Delta_{S_L}(\hat{r}(x,t) - r_T(x)) && x \in \mathcal{N}_{\partial S_F} && \text{(12c)}
\end{aligned}
$$

where $r_T(x),\ x \in \mathcal{N}_{\partial S_F}$ denotes the desired target position of leader x, and where we use Δ_{S_L} to denote the Laplacian operator defined solely over the subgraph S_L, i.e.

$$\Delta_{S_L} f(x) \doteq - \sum_{y \sim x,\ y \in \mathcal{N}_{S_L}} \partial_y^2 f(x).$$

Now, under the assumption that S_L is connected, and, by exactly the same reasoning as for the standard rendez-vous problem, the leaders will converge to positions $r_L(x)$ such that $\partial_y r_L(x) = \partial_y r_T(x),\ \forall x, y \in \mathcal{N}_{S_L}$. In other words, no convergence to a predefined point is achieved. Rather, this control law ensures that the leaders arrive at a translationally invariant target formation.

Note that the details of the leaders' motion is not crucial and this particular choice is but one of many possibilities. However, this choice is appealing in that it makes the information flow explicit, and the leaders only need access to the positions (and target locations) of their neighboring leaders in order to compute their motion. As such the decentralized character of the algorithm is maintained.

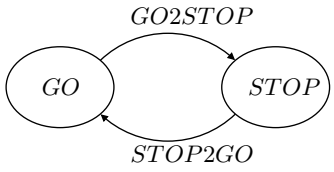

Fig. 3. The hybrid automaton implementing the Stop-Go policy

In order for fully specify the hybrid Stop-Go leader policy depicted in Fig. 3, transition rules are needed as well. As before, let Ω_L denote the leader-polytope and let $d(\mu, \Omega_L)$ denote the signed distance

$$d(\mu, \Omega_L) \doteq \zeta_{\Omega_L}(\mu) \min_{x \in \partial\Omega_L} \|\mu - x\|_2,$$

where $\|\cdot\|_2$ denotes the Euclidean 2-norm, and where $\zeta_{\Omega_L}(\mu) = -1$ if $\mu \in \Omega_L$ and $+1$ otherwise. Using this distance measure we let the two guards be given by

$$GUARD_{GO2STOP}: \ \exists y \in \mathcal{N}_{S_F} \mid d(r(y,t), \Omega_L) \geq 0? \tag{13}$$

$$GUARD_{STOP2GO}: \ d(r(y,t), \Omega_L) < -\epsilon \ \forall y \in \mathcal{N}_{S_F}? \tag{14}$$

Note that the guard $STOP2GO$ is crossed only if the following assumptions are verified:

Assumption 1. *Let $\hat{h}(\cdot,t)$ be the solution to (5) for $\bar{r}(\cdot) = \hat{r}(\cdot,t)$, $\forall t \geq 0$ and consider the set $\Omega_L^\epsilon(t) = \{y \in \Omega_L(t) : \text{dist}(y, \partial\Omega_L(t)) > \epsilon\}$. Then*

1. *$\Omega_L^\epsilon(t)$ is nonempty, $\forall t \geq 0$;*
2. *$\text{Co}(\{\hat{h}(x,t), x \in \mathcal{N}_{S_F}\}) \subset \Omega_L^\epsilon(t)$.*

In particular, Assumption 1 implies that Ω_L must be full-dimensional at all times and "sufficiently fat" along every direction. Conditions relating property 2 of Assumption 1 to the graph topology are currently under investigation. A few comments must be made about the computation and communication requirements that these guards give rise to. If two leaders are located at the end-points of the same face of Ω_L, then they must be able to determine if any of the followers are in fact on this face. This can be achieved through a number of range sensing devices, such as ultrasonic, infra-red, or laser-based range-sensors. Moreover, in order for all leaders to transition between modes in unison, they must communicate between them, which means that either ∂S_F is a complete graph, or that multi-hop strategies are needed. In either way, a minimal requirement for these mode transitions to be able to occur synchronously, without having to rely on information flow across follower-agents, is that ∂S_F must be connected.

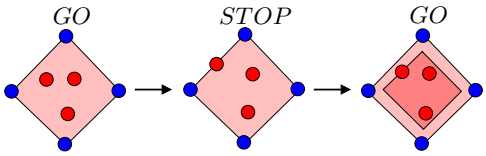

Fig. 4. A hysteresis-based transition strategy avoids Zeno executions

The hysteresis threshold $\epsilon > 0$ in the $STOP2GO$ guard (see Fig. 4) is needed in order to avoid Zeno executions, as seen from the following argument: The distance between any points in Ω_L is less than the diameter of Ω_L. We let ρ_{Ω_L} denote the supremum of these diameters during an execution, and note that since the leaders are under our control, ρ_{Ω_L} can be prevented from being unbounded, and we state this as an assumption:

Assumption 2. *$\exists M < \infty$ such that $\rho_{\Omega_L} \leq M$.*

Note that the Laplacian control law used for controlling the leaders is but one of many possible control strategies. As such, we can always use for example a plan-based leader control law if the Laplacian control law was to violate the assumption.

Under the above-mentioned assumption we have

$$\|\dot{r}(x,t)\| = \|\Delta r(x,t)\| \leq \sum_{y\sim x} \|\partial_y r(x)\| \leq \sum_{y\sim x} \rho_{\Omega_L} \leq N\rho_{\Omega_L}, \quad \forall x \in \mathcal{N}_{S_F}.$$

Now, in order for the system to leave the Stop mode, at least one follower agent must have traveled at least a distance ϵ, which in turn implies that the system will always stay for a time greater than or equal to $\epsilon/N\rho_{\Omega_L}$ in the Stop mode. And, in order for the system to exhibit Zeno executions, a necessary condition is that the difference between the transition times must approach zero [18]. And, since this is not the case here, the non-Zeno property is established.

In the following section, an example is given that describes the operation and the feasibility of the proposed Stop-Go control policy.

5 Examples

The previous sections show that the polytope spanned by the leaders is invariant to the followers and the hybrid control strategy is non-Zeno. In this section, an example is given to show the validity of the proposed control method.

A scenario where three leaders (black) maneuver four followers (white) is investigated here. The initial position and final position of the leaders are $r(x,0) = \{(1,-3),(0,-1),(0,1)\}$ and $r_T(x) = \{(0,-2),(1,2),(2,-2)\}$ respectively. The followers are indexed from 1 to 4 and the leaders from 5 to 7.

During the maneuvering, the Stop-Go policy is adopted, i.e. the followers are governed by the Laplacian control, while the leaders dynamics are only affected by other leaders, as in (12c).

In Fig. 5, the snap-shots of the herding process are shown. The magnitude of the velocities of the agents are shown in Fig. 6, where we can see the instances when the leaders stop to make sure that the followers remain inside the leader-polytope. The snap-shots of the transition instances are shown in Fig. 7.

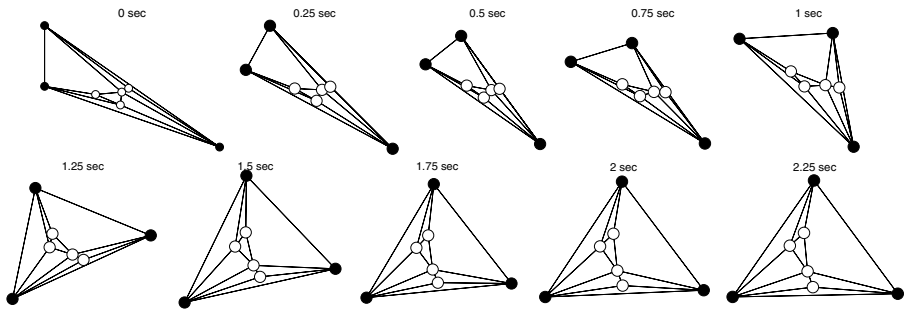

Fig. 5. A herding process where 4 followers (white) are herded by 3 leaders (black), who use the hybrid Stop-Go control policy

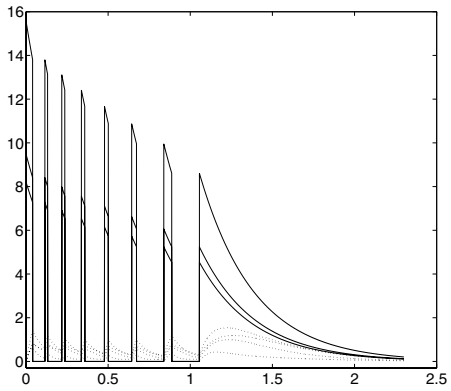

Fig. 6. The magnitude of the velocities of the agents in the Stop-Go herding process. Solid lines correspond the velocity of the leaders while dashed lines correspond to those of the followers.

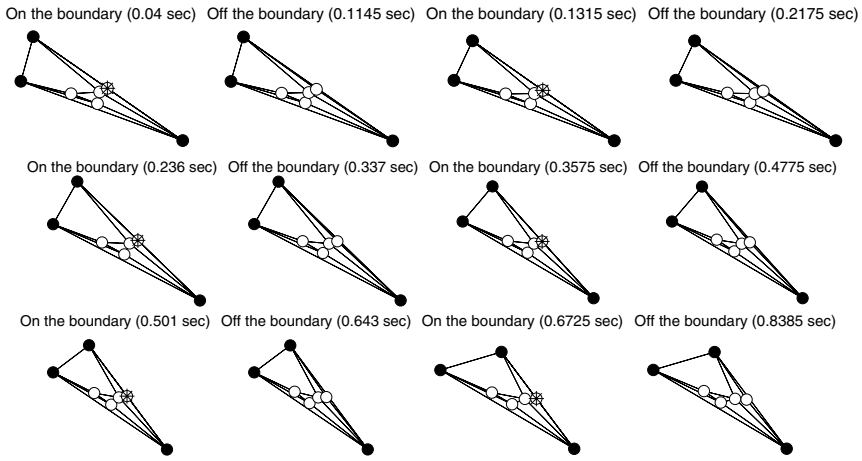

Fig. 7. Time instances when transitions occur. (The asterisk denotes the particular follower who touches the boundary.)

6 Extensions: Liveness Issues and Hierarchical Control

As already mentioned, the proposed solution is non-Zeno. However, as it is currently defined, the Stop-Go policy may be blocking in the sense that the system never leaves the Stop mode. One remedy to this problem is to allow the containment to be slightly less tight. In other words, we can select different guards, e.g.

$$GUARD_{GO2STOP}: \ \exists y \in \mathcal{N}_{S_F} \mid d(r(t,y), \Omega_L) > \delta? \tag{15}$$

$$GUARD_{STOP2GO}: \ d(r(t,y), \Omega_L) \leq 0 \ \forall y \in \mathcal{N}_{S_F}? \tag{16}$$

What this means is that we do not enter the Stop mode until a follower is $\delta > 0$ outside Ω_L.

Assumption 3. *The set Ω_L is full dimensional at all times and* $\mathrm{Co}(\{\hat{r}_T(x), x \in \mathcal{N}_{\partial S_F}\})$ *is full dimensional.*

Under Assumption 3, the size of Ω_L is lower bounded at all times by a positive constant and hence, by virtue of Theorem 2, every follower will eventually get back in Ω_L in finite time (since the leaders are stationary in the Stop mode). This argument proves that the system is live in the sense of always leaving the Stop mode eventually.

However, liveness is not enough. We moreover must ensure that we do in fact reach the target location.

Under Assumption 2, it holds $\|\dot{r}(x,t)\| \leq N(\rho_{\Omega_L} + \delta)$ and we can repeat the non-Zeno argument from Section 4 in order to see that the system always stays in the Go mode for a time greater than or equal to $\delta/(N(\rho_{\Omega_L} + \delta))$. In fact, this bound can be made tighter by virtue of Theorem 3, since we do not need to take the followers on $\partial\Omega_L$ into account because the motion of the followers is such that their velocities will never point away from Ω_L. In other words, a transition from Go to Stop occurs when the leaders "catch up" with the followers rather than when the followers move away from Ω_L. As a result, in a non-blocking system the leaders will be given infinitely many opportunities to move during a finite (bounded away from zero) time horizon, which implies convergence to the target location as long as the the leaders would in fact end up at the target location under the influence of the Go mode alone without the leader polytope degenerating to a convex polytope of a reduced dimension.

Another direction in which additional improvements can be expected is to reduce the necessary information flow by imposing a hierarchical structure on the formation. This can for instance be achieved by organizing the agents into M layers such that each agent in layer i, $i = 1, 2, \ldots, M-1$, is a follower of the agents in its upper layer, $i+1$, i.e. $\mathcal{N}_{\partial S_{i+1}} \subseteq \mathcal{N}_{S_i} \cup \mathcal{N}_{S_{i+2}}\ \forall\, i = 1, \cdots, M-2$, where S_i denotes the subgraph corresponding to layer i and ∂S_{i-1} is the (non-empty) boundary of S_{i-1} in the host graph G. In such a setting, the Stop-Go control policy would still be applicable. Only the agents of the outermost layer would be given target locations, $r_T(x)$. All other layers simply obey the Laplacian control strategy unless that layer's boundary is intersected by an agent belonging to layer $i-1$, at which point they would halt their motion. This, however, is a research direction that is left to the future, and we simply mention it here as a possible and promising extension.

7 Conclusions

In this paper we present a hybrid Stop-Go control policy for the leaders in a multi-agent containment scenario. In particular, the control strategy allows us to transport a collection of follower-agents to a target area while ensuring that they stay in the convex polytope spanned by the leaders. The enabling results

needed in order to achieve this is that, for stationary leaders, the followers in a connected interaction graph will always converge to locations in the leader-polytope. Additional extensions to the proposed control strategy are given in order to ensure certain liveness properties and we outline how the proposed methods lend themselves easily to generalizations to hierarchical information exchange strategies. Examples are moreover presented in order to stress the viability of the proposed approach.

Acknowledgment

The work by Giancarlo Ferrari-Trecate was partially supported by the European Commission under the Network of Excellence HYCON, contract number FP6-IST-511368. The work by Magnus Egerstedt and Meng Ji was supported by the U.S. Army Research Office through Grant #99838.

References

1. Parker, L.E., Kannan, B., Fu, X., Tang, Y.: Heterogeneous mobile sensor net deployment using robot herding and line-of-sight formations. In: Proceedings of IEEE International Conference on Intelligent Robots and Systems. (2003) 681–689
2. Vaughan, R., Sumpter, N., Henderson, J., Frost, A., Cameron, S.: Experiments in automatic flock control. Journal of Robotics and Autonomous Systems **31** (2000) 109–117
3. Desai, J., Ostrowski, J.P., Kumar, V.: Controlling formations of multiple mobile robots. In: Proc. IEEE Int. Conf. Robot. Automat. (1998) 2864–2869
4. Ji, M., Muhammad, A., Egerstedt, M.: Leader-based multi-agent coordination: Controllability and optimal control. (Submitted to the American Control Conference, Minneapolis, MN, June 2006)
5. Tanner, H., Pappas, G., Kumar, V.: Leader to formation stability. IEEE Transactions on Robotics and Automation **20**(3) (2004) 443–455
6. Tanner, H., Jadbabaie, A., Pappas, G.: Flocking in fixed and switching networks. (2005) `http://www.seas.upenn.edu/~pappasg/publications.html#journals`. Submitted.
7. Saber, R.O.: A unified analytical look at Reynolds flocking rules. Technical Report CIT-CDS 03-014, California Institute of Technology (2003)
8. Muhammad, A., Egerstedt, M.: Connectivity graphs as models of local interactions. (To appear in Journal of Applied Mathematics and Computation, 2005)
9. Egerstedt, M., Martin, C.: Conflict resolution for autonomous vehicles: A case study in hierarchical control design. International Journal of Hybrid Systems **2**(3) (2002) 221–234
10. Sussmann, H.: A maximum principle for hybrid optimal control problems. 38th IEEE Conference on Decision and Control (1999)
11. Ferrari-Trecate, G., Buffa, A., Gati, M.: Analysis of coordination in multiple agents formations through Partial difference Equations. Technical Report 5-PV, IMATI-CNR (2004) `http://www-rocq.inria.fr/who/Giancarlo.Ferrari-Trecate/publications.html`.
12. Bollobás, B.: Modern graph theory. Graduate texts in Mathematics. Springer-Verlag (1998)

13. Bensoussan, A., Menaldi, J.L.: Difference equations on weighted graphs. Journal of Convex Analysis (Special issue in honor of Claude Lemaréchal) **12**(1) (2005) 13–44
14. Ferrari-Trecate, G., Buffa, A., Gati, M.: Analysis of coordination in multi-agent systems through partial difference equations. Part I: The Laplacian control. 16th IFAC World Congress on Automatic Control (2005)
15. Jadbabaie, A., Lin, J., Morse, A.S.: Coordination of groups of mobile autonomous agents using nearest neighbor rules. IEEE Trans. on Automatic Control **48**(6) (2003) 988 – 1001
16. Olfati-Saber, R., Murray, R.: Consensus problems in networks of agents with switching topology and time-delays. IEEE Trans on Autom. Control **49**(9) (2004) 101–115
17. Dautray, R., Lions, J.L.: Mathematical analysis and numerical methods for science and technology. Vol. 5-6: Evolution problems I-II. Springer-Verlag, Berlin (1992)
18. Johansson, K., Egerstedt, M., Lygeros, J., Sastry., S.: Regularization of zeno hybrid automata. Systems and Control Letters **38** (1999) 141–150

Appendix A

This Appendix is devoted to the proof of Theorem 2. We start by introducing a basic result on polytopes.

Lemma 1. *Consider the polytope* $P = \mathrm{Co}(X)$ *where* $X = \{x_i \in \mathbb{R}^d : i = 1, \ldots, L\}$ *and let* X_1 *be a proper subset of* X*. If* $x \in \mathrm{Co}(X_1)$, $\forall x \in X \backslash X_1$, *then* $P = \mathrm{Co}(X_1)$.

Proof: The conditions $x \in X \backslash X_1$ and $x \in \mathrm{Co}(X_1)$ imply that x is not a vertex of P. Then, X_1 includes all vertexes of P, thus proving that $P = \mathrm{Co}(X_1)$. ■

Lemma 2. *Let* G *be a host graph,* S *a subgraph of* G *and* T_1 *a proper subgraph of* S *and* ∂T_1 *the boundary of* T_1 *in* G*. Consider* $\bar{x} \in \mathcal{N}_{\partial T_1} \cap \mathcal{N}_S$*, and let* $r \in L^2(G)$ *be a function verifying*

$$r(\bar{x}) \in \mathrm{Co}(\{r(y) : y \in \mathcal{P}(\bar{x})\}) \tag{17}$$

$$r(x) \in \mathrm{Co}(\{r(y) : y \in \mathcal{N}_{\partial T_1}\}), \quad \forall x \in \mathcal{N}_{T_1} \tag{18}$$

Let T_2 *be the subgraph associated with* $\mathcal{N}_{T_1} \cup \{\bar{x}\}$ *and* ∂T_2 *be the boundary of* T_2 *in* G*. Then, for all* $x' \in \mathcal{N}_{T_2}$ *it holds*

$$r(x') \in \mathrm{Co}(\{r(y) : y \in \mathcal{N}_{\partial T_2}\}). \tag{19}$$

Proof: From (17), one has that all $\bar{r} \in \{r(y) : y \in \mathcal{N}_{T_2}\}$ verify $\bar{r} \in P$ where $P = \mathrm{Co}(\{r(y) : y \in \mathcal{N}_{\partial T_1} \cup \mathcal{P}(\bar{x})\})$. In particular, if $x \in \mathcal{P}(\bar{x}) \cap T_1$ one has that $r(x) \in \mathrm{Co}(\{r(y) : y \in \mathcal{N}_{\partial T_1}\})$. Recalling (18) and that $\bar{x} \in \mathcal{N}_{\partial T_1}$ one can apply Lemma 1 and obtain

$$P = \mathrm{Co}(\{r(y) : y \in (\mathcal{N}_{\partial T_1} \backslash \{\bar{x}\}) \cup (\mathcal{P}(\bar{x}) \backslash \mathcal{N}_{T_1})\}).$$

The proof is concluded by realizing that

$$\mathcal{N}_{\partial T_2} = (\mathcal{N}_{\partial T_1} \backslash \{\bar{x}\}) \cup (\mathcal{P}(\bar{x}) \backslash \mathcal{N}_{T_1}).$$

■

We are now in a position to prove Theorem 2.

Proof: (Theorem 2.)

Let $p = x_0 x_1 \dots x_L$ be a path going through all nodes of S_1. Since $\Delta h(x) = 0$, $\forall x \in \mathcal{N}_{S_F}$ from (1) one has

$$h(x) = \frac{1}{|\mathcal{P}(x)|} \sum_{y \in \mathcal{P}(x)} h(y), \ \forall x \in \mathcal{N}_{S_F}$$

which implies that

$$h(x) \in \mathrm{Co}\{h(y) : y \in \mathcal{P}(x)\} \tag{20}$$

We will prove the theorem using a recursive argument on the nodes composing $\mathcal{P}$. First, note that $x_1 \in \mathcal{P}(x_0)$. Let T_1 and T_2 be the subgraphs associated with $\{x_0\}$ and $\mathcal{N}_{S_1} \cup \{x_1\}$, respectively. Lemma (2) can be applied with $\bar{x} = x_1$. Indeed, (17) amounts to (20) for $x = x_0$ and (18) amounts to (20) for $x = x_1$. Then, from (19) we have

$$h(x) \in \mathrm{Co}(\{h(y) : y \in \mathcal{N}_{\partial T_2}\}), \ \forall x \in \mathcal{N}_{T_2}$$

Now, we denote by $S^{(i)}$, $i < L$ the subgraph of S_1 associated with the $i+1$ nodes $\{x_0, x_1, \dots, x_i\}$ and by $\partial S^{(i)}$ its boundary in G. Assume now that at the i-th step, $i < L$ we have

$$h(x) \in \mathrm{Co}(\{h(y) : y \in \mathcal{N}_{\partial S^{(i)}}\}), \ \forall x \in \mathcal{N}_{S^{(i)}} \tag{21}$$

We need to prove that:

$$h(x) \in \mathrm{Co}(\{h(y) : y \in \mathcal{N}_{\partial S^{(i+1)}}\}), \ \forall x \in \mathcal{N}_{S^{(i+1)}}. \tag{22}$$

Note that $x_{i+1} \in \mathcal{P}(x_i)$. Set $T_1 = S^{(i)}$ and let T_2 be the graph associated with $\mathcal{N}_{S^{(i)}} \cup \{x_{i+1}\}$. Lemma (2) can be applied with $\bar{x} = x_{i+1}$. Indeed, (17) amounts to (20) for $\bar{x} = x_{i+1}$ and (18) amounts to (21). Then, from (19) we have

$$h(x) \in \mathrm{Co}(\{h(y) : y \in \mathcal{N}_{\partial T_2}\}), \ \forall x \in \mathcal{N}_{T_2}.$$

Since, $T_2 = S^{(i+1)}$, formula (8) is proved. If S is connected, the result holds also for $S = S_1$. If S is not connected, we apply (8) on each connected component S_i, $1 \le i \le n$, and, by simple algebra, obtain

$$h(x) \in \mathrm{Co}(\{h(y), \ y \in \mathcal{N}_{\partial S_1} \cup \mathcal{N}_{\partial S_2} \cup \dots \cup \mathcal{N}_{\partial S_n}\}).$$

The proof that each follower lies in the convex hull of the leaders positions is ended by realizing that $\mathcal{N}_{\partial S} = \mathcal{N}_{\partial S_1} \cup \mathcal{N}_{\partial S_2} \cup \dots \cup \mathcal{N}_{\partial S_n}$.

The fact that the full dimensionality of Ω_L implies that $h(x) \in \Omega_L \backslash \partial\Omega_L$, $\forall x \in \mathcal{N}_{S_F}$ is proved by contradiction. Let $x \in \mathcal{N}_{S_F}$ be such that $h(x) \in \partial\Omega_L$ and denote with χ the supporting hyperplane of Ω_L such that $h(x) \in \Omega_L \cap \chi$. Then, since $h(x) \in \mathrm{Co}(\{h(y), \ y \in \mathcal{P}(x)\})$, all $y \in \mathcal{P}(x)$ verify $h(y) \in \partial\Omega_L \cap \chi$. Iterating the argument over the followers lying on χ, one would find that also all leaders $x \in \mathcal{N}_{\partial S_F}$ lie on χ and this contradicts the fact that Ω_L is full dimensional. ■