Proceedings of the
47th IEEE Conference on Decision and Control
Cancun, Mexico, Dec. 9-11, 2008

Stable Emergent Agent Distributions under Sensing and Travel Delays

Jorge Finke, Brandon J. Moore, and Kevin M. Passino
Department of Electrical and Computer Engineering
The Ohio State University
{finkej, mooreb, passino}@ece.osu.edu

Abstract— **In order for a team of cooperating agents to achieve a group goal (such as searching for targets, monitoring an environment, etc.) those agents must be able to share information and achieve some level of coordination. Since realistic methods of communication between agents have limited range and speed, the agents' decision-making strategies must operate with incomplete and outdated information. Moreover, in many situations the agents must travel to particular locations in order to perform various tasks, and so there will also be a delay between any particular decision and its effect. In this paper we develop an asynchronous framework that models the behavior of a group of agents that is spatially distributed across a predefined area of interest. We derive general conditions under which the group is guaranteed to converge to a specific distribution within the environment without any form of central control and despite unknown but bounded delays in sensing and travel. The achieved distribution is optimal in the sense that the proportion of agents allocated over each area matches the relative importance of that area. Finally, based on the derived conditions, we design a cooperative control scheme for a multi-agent surveillance problem. Via Monte Carlo simulations we show how sensing and travel delays and the degree of cooperation between agents affect the rate at which they achieve the desired coverage of the region under surveillance.**

I. Introduction

In systems made of a large number of autonomous self-driven components (called agents) "cooperation" describes the process of working together in order to meet a common group objective. Cooperation inherently requires that agents share information, e.g., either through direct agent-to-agent communication, indirectly with the aid of intermediate agents, or via visual cues in the environment. Through communication each individual agent develops a perception about both the state of other agents and the environment. In order to control a cooperative system it is necessary to design distributed decision-making strategies which lead to a desired group objective despite each individual agent having limited and possibly inaccurate information. Not surprisingly, researchers have devoted plenty of attention to techniques from the field of distributed algorithms and computation, where information flow constraints considerably impact the performance of distributed processors [1], thereby creating challenges similar to those in cooperative systems.

For example, in iterative computing, agreement algorithms are widely used to reconcile the updates made by the individual processors in a distributed control scheme where several processors update the same set of variables (e.g., see Section 7.7 in [1]). Similarly, in cooperative systems, agents must often decide on a particular variable of interest and rely on agreement algorithms to achieve coordination of the group [2]-[6] (e.g., to agree to a heading and speed for movement in a formation). While many such models have been developed under various information flow constraints (see [7], [8] for a recent survey on agreement problems in multi-agent coordination), they usually require a central assumption in their results, namely that the updated state of each agent is a strict convex combination of its own current state and the current or past states of the agents to which it connects. Although limited in more general contexts, such as when the state of the system evolves in partially obstructed Euclidean spaces [9], the convexity assumption offers a practical mathematical tool which is often used to guarantee the desired behavior of a group as a whole.

Broadly speaking, the agent strategies we introduce here resemble agreement algorithms in the sense that to achieve a common group objective, agents must divide themselves into a fixed number of sub-groups (each of which may represent either a portion of the total environment or a different type of task) while consenting on "gains" which are associated to every subgroup. More specifically, we refer to these sub-groups as nodes (because of a connection to the terminology of graph theory we will see later) and let the number of agents at a node represent the state of the node. Assuming an inverse relationship between the state of a node and its associated gain, we study how the agents' motion dynamics across all nodes may lead to a desired agent distribution resulting in equal gains over all nodes. Our framework contemplates three aspects of the system in particular. First, there exists a mapping between the state space (i.e., the simplex representing all possible distributions of the agents across the nodes) and the space of possible gains that results from all such distributions. We derive *general* sensing and motion conditions which identify key requirements on the amount of agents that may (or must sometimes) leave certain nodes to guarantee convergence to the desired distribution, while relying only on the individual agents' outdated perceptions about the gains of the few nodes they can sense. Second, depending on the *total* number of agents in the system, there may exist so-called truncated nodes which can never reach the desired gain regardless of their state (i.e., if the gain associated with a node remains relatively low even in the presence of only a few agents, while the distribution of all other agents does not degrade the gains of all other nodes to a similar level). Analyzing this possibility is important because the

978-1-4244-3124-3/08/$25.00 ©2008 IEEE

existence of truncated nodes may significantly influence the group's distribution (and thus the resulting gains) depending on the particular scenario. And finally, since our approach is based on techniques used in diffusion algorithms for load balancing (where loads move from heavily loaded processors to lightly loaded neighbor processors [10]-[12]), there is no requirement to form convex combinations of the current or past states of the nodes. Although some general ideas are similar to agreement algorithms, the convergence of diffusion algorithms cannot be derived from the corresponding results for agreement algorithms [2]-[6].

The remainder of this paper is organized as follows. In Section II we define a basic mathematical formulation of our problem and present a class of distributed control algorithms that will solve it. Our results in Section III extend the load balancing theory in [12], [13] by taking into account: (i) that the "virtual load" is a *nonlinear* function of the state, and (ii) the presence of sensing and travel delays. Our analytical results show that although delays may increase the time for the agents to converge, a global distribution pattern will still be reached, provided that agents at any node have a perception (possibly outdated, but by a finite delay) about the gains of any other node, i.e., when a fully connected topology represents the information flow structure between the nodes. Since fully connected topologies are rarely applicable we present similar results that show that under stronger conditions on the total number of available agents, the desired distribution can still be achieved for a general topology under only minimal restrictions on the graph topology.

Finally, in Section IV we apply the theory presented here to design cooperative control strategies for multi-agent surveillance problems. We extend our previous results in [14], [15] by quantifying the degree of cooperation between agents (i.e., the willingness of the agents of working together in order to meet the group objective) and show how sensing and travel delays considerably impact the degree to which agents should cooperate so that (on average) they achieve the desired distribution as fast as possible.

II. The model

Assume that there are N nodes, each of which is characterized by an associated gain. Define the gain function of node i as $s_i(x_i)$, where $x_i \in \mathbb{R},\ x_i \geq \varepsilon_p$ is a scalar that represents the amount of agents located at node $i \in \mathcal{H}$, $\mathcal{H} = \{1, ..., N\}$, and $\varepsilon_p \geq 0$ is the minimum amount of agents required at any node. In some cases, if for example, $s_i(x_i) = 1/x_i$, then we require that $\varepsilon_p > 0$ so that the gain function of node i is well-defined at any state. In other cases, if for example $s_i(x_i) = e^{-x_i}$, then we may let $\varepsilon_p > 0$ solely to enforce that a certain amount of agents always remain at any node. Furthermore, assume the following:

- *A fixed group size:* Let $\sum_{i=1}^{N} x_i = P$, where $P > N\varepsilon_p$ is a constant so there is a fixed amount of agents distributed across all nodes.
- *Positive gains:* The gain functions $s_i(x_i) > 0$ for all $i \in \mathcal{H}$, and all $x_i \in [\varepsilon_p, P]$.
- *Gain changes are related to changes in the amount of agents:* For all gain functions $s_i(x_i)$, $i \in \mathcal{H}$, there exists two constants $a_i \geq b_i > 0$ such that
$$-a_i \leq \frac{s_i(y_i) - s_i(z_i)}{y_i - z_i} \leq -b_i \qquad (1)$$
for any $y_i, z_i \in [\varepsilon_p, P]$, $y_i \neq z_i$. Note that Eq. (1) implies that the gain associated with each node decreases with an increasing amount of agents at that node, and eliminates the possibility that a very small difference in agents may result in an unbounded change in gain.

To model interconnections between nodes we consider a general graph described by $\mathcal{G}(\mathcal{H}, \mathcal{A})$ with topology $\mathcal{A} \subset \mathcal{H} \times \mathcal{H}$. Let $\mathcal{N}(i) = \{j : \exists (i,j) \in \mathcal{A}\}$ denote the *neighboring* nodes of node i, i.e., the nodes where agents at node i can move to and whose gain they can sense. If $(i,j) \in \mathcal{A}$ then $i \neq j$ and $(j,i) \in \mathcal{A}$, which means that agents at node i can sense the gain of node j and can move from node i to node j, and vice versa (i.e., if an agent is at node i and can move to node j (sense the gain at node j), then agents at node j can also move from node j to node i (sense the gain at node i, respectively)). For all $x_i(k)$, agents at node i at time k "sensing node j" means that they have a *perception* $p_j^i(k)$ of the gain of node j, which is a delayed value of the current gain of node j. In particular, if $(i,j) \in \mathcal{A}$, then we assume that any changes in the gain of node j at time k can be sensed by agents at node i by time $k + B_s - 1$ for some $B_s > 0$. In other words, we assume that there exists a constant B_s such that $p_j^i(k) \in \{s_j(x_j(k')) : k - B_s < k' \leq k\}$. Likewise, for some agents at node i "moving to node j" at time k implies that they start traveling away from node i at time k and will arrive at node j at some time k', for $k < k' \leq k + B_t - 1$ and some constant $B_t > 1$ (note that the maximum travel delay is $B_t - 1$). We assume that agents at node i know the value of $s_i(x_i(k))$ so that $p_i^i(k) = s_i(x_i(k))$, and are assumed to know $x_i(k)$. We also assume that for every $i \in \mathcal{H}$, there must exist some $j \in \mathcal{H}$ such that $j \in \mathcal{N}(i)$ and that there exists a path between any two nodes (which ensures that every node is connected to the graph $\mathcal{G}(\mathcal{H}, \mathcal{A})$).

Let $\mathcal{X} = \mathbb{R}^{N \times (B_s + NB_t)}$ define the set of states. Every state $x(k) \in \mathcal{X}$ is composed of (i) the total amount of agents *located at the nodes* for all k' such that $k - B_s < k' \leq k$ (which we will capture by $x_n(k)$); and (ii) the total amount of agent *traveling between nodes* for all k' such that $k - B_t < k' \leq k$ (which we will capture by $x_t(k)$). In particular, $x_n(k) \in \mathbb{R}^{N \times B_s}$ and $x_t(k) \in \mathbb{R}^{N \times NB_t}$ are defined as,
$$\begin{aligned} x_n(k) &= \begin{bmatrix} x_1(k) & \dots & x_1(k - B_s + 1) \\ \vdots & \ddots & \vdots \\ x_N(k) & \dots & x_N(k - B_s + 1) \end{bmatrix} \\ x_t(k) &= \begin{bmatrix} x_t^1(k) & \dots & x_t^N(k) \end{bmatrix}, \end{aligned}$$
where
$$x_t^i(k) = \begin{bmatrix} x_{i\to 1}(k) & \dots & x_{i\to 1}(k - B_t + 1) \\ \vdots & \ddots & \vdots \\ x_{i\to N}(k) & \dots & x_{i\to N}(k - B_t + 1) \end{bmatrix}$$

and $x_{i\to j}(k)$ denotes the amount of agents that are traveling from node i to node j at time k. The state of the system is defined as $x(k) = [x_n(k),\ x_t(k)]$. Next, we want to define a set of states, such that any state $x(k)$ that belongs to this set exhibits the following desired properties:

Property 1: Agents at time k are distributed such that:

- All nodes with more than the minimum amount of agents ε_p have equal gains;
- Any node that does not have the same gain as its neighboring nodes must have a lower gain and the minimum amount of agents ε_p only.

Property 2: There are no agents traveling between any two nodes at times k, $k-1$, …, $k - B_t + 1$.

Property 3: At time k every agent at any node $i \in \mathcal{H}$ has an accurate perception about the gain of its neighboring nodes $\mathcal{N}(i)$ (i.e., it will sense the actual gain at time k).

Let $(\cdot)_{ij}$ denote the element in row i and column j of its matrix argument. Let $\mathcal{S} = \{1, \ldots, B_s\}$, and $\mathcal{T} = \{1, \ldots, B_t\}$. Note that any distribution of agents such that the state belongs to the set

$$\begin{aligned} \mathcal{X}_b \;=\; & \{x \in \mathcal{X} : \forall i, p \in \mathcal{H},\ q \in \mathcal{T}, \\ & (x_t^i)_{pq} = 0;\ \ \forall i \in \mathcal{H},\ j \in \mathcal{S}, \\ & \textit{either}\ \ s_i((x_n)_{ij}) = s_p((x_n)_{pq}), \\ & \qquad \forall p \in \mathcal{H},\ q \in \mathcal{S} \text{ such that } (x_n)_{pq} \neq \varepsilon_p, \\ & \qquad \text{and } s_i((x_n)_{ij}) \geq s_p((x_n)_{pq}), \\ & \qquad \forall p \in \mathcal{H},\ q \in \mathcal{S} \text{ such that } (x_n)_{pq} = \varepsilon_p; \\ & \textit{or}\ \ (x_n)_{ij} = \varepsilon_p\} \end{aligned} \tag{2}$$

possesses the desired properties. In particular, note that since the gain of all nodes has been fixed since time $k - B_s + 1$, we are guaranteed agents at every node have an accurate perception of all of its neighboring nodes. Next, we specify the set of events that capture the agents dynamics across the nodes, and define sensing and motion conditions that ensure that the desired agent distribution will be reached.

Let $e_{\alpha(i,k)}^{i\to\mathcal{N}(i)}$ represent the event that some agents from node $i \in \mathcal{H}$ start moving to neighboring nodes $\mathcal{N}(i)$ at time k, where $\alpha(i,k)$ is a list $(\alpha_j(i,k), \alpha_{j'}(i,k), \ldots, \alpha_{j''}(i,k))$ such that $j < j' < \ldots < j''$ and $j, j', \ldots, j'' \in \mathcal{N}(i)$ whose elements $\alpha_j(i,k) \geq 0$ denote the amount of agents that start moving from node i to node $j \in \mathcal{N}(i)$ at time k (note that the the size of $\alpha(i,k)$ is $|\mathcal{N}(i)|$). For convenience, we denote this list by $\alpha(i,k) = (\alpha_j(i,k) : j \in \mathcal{N}(i))$. Let $\left\{e_{\alpha(i,k)}^{i\to\mathcal{N}(i)}\right\}$ represent the set of all possible combinations of how agents can move from node i to neighboring nodes $\mathcal{N}(i)$ at any time k. Furthermore, let $e_{\beta(i,k)}^{i\leftarrow\mathcal{N}(i)}$ represent the event that agents from some neighboring nodes arrive at node i, where $\beta(i,k) = (\beta_j(i,k), \beta_{j'}(i,k), \ldots, \beta_{j''}(i,k))$ such that $j < j' < \ldots < j''$ and $j, j', \ldots, j'' \in \mathcal{N}(i)$ is a list composed of elements $\beta_j(i,k)$ that denote the amount of agents that arrive from a neighboring node $j \in \mathcal{N}(i)$ at node i at time k. Again, for convenience we denote this list by $\beta(i,k) = (\beta_j(i,k) : j \in \mathcal{N}(i))$. Let $\left\{e_{\beta(i,k)}^{i\leftarrow\mathcal{N}(i)}\right\}$ denote the set of all possible arrivals at node i at any time k. Finally, let the set of events be described by

$$\mathcal{E} = \left\{\mathcal{P}\left(\left\{e_{\alpha(i,k)}^{i\to\mathcal{N}(i)}\right\}\right) \bigcup \mathcal{P}\left(\left\{e_{\beta(i,k)}^{i\leftarrow\mathcal{N}(i)}\right\}\right)\right\} - \{\emptyset\}$$

where $\mathcal{P}(\cdot)$ denotes the power set of its argument. Each event $e(k) \in \mathcal{E}$ is defined as a set, with each element of $e(k)$ representing the departure of agents from $i \in \mathcal{H}$ or the arrival of agents at node $i \in \mathcal{H}$, and multiple elements in $e(k)$ representing the simultaneous movements of agents, i.e., agent departures and arrivals to multiple nodes.

An event $e(k) \in \mathcal{E}$ may only occur if it is in the set defined by an "enable function," denoted by $g : \mathcal{X} \to \mathcal{P}(\mathcal{E}) - \{\emptyset\}$. State transitions are defined by the operators $f_e : \mathcal{X} \to \mathcal{X}$, where $e \in \mathcal{E}$. By specifying g and f_e for $e(k) \in g(x(k))$ we define the agents' sensing and motion conditions:

- Event $e(k) \in g(x(k))$ if (a), (b), and (c) below hold:

  (a) For all $e_{\alpha(i,k)}^{i\to\mathcal{N}(i)} \in e(k)$, it is the case that:

$$\begin{aligned} (i) \qquad & \alpha_m(i,k) = 0 \ \text{ if } \ s_i(x_i(k)) \geq p_m^i(k) \\ (ii) \qquad & x_i(k) - \sum_{m\in\mathcal{N}(i)} \alpha_m(i,k) \geq \varepsilon_p \\ (iii) \qquad & s_i(x_i(k)) + a_i \sum_{m\in\mathcal{N}(i)} \alpha_m(i,k) \\ & \leq p_{j^*}^i(k) - a_{j^*}\alpha_{j^*}(i,k) \\ (iv) \qquad & \alpha_{j^*}(i,k) \geq \frac{\gamma_{ij^*}}{b_{j^*}} \left[p_{j^*}^i(k) - s_i(x_i(k))\right], \\ & \text{if } x_i(k) \geq \alpha_{j^*}(i,k) + \varepsilon_p \ \text{ and} \\ & \alpha_{j^*}(i,k) = x_i(k) - \varepsilon_p, \ \text{otherwise} \end{aligned}$$

  where $j^* \in \left\{j : p_j^i(k) \geq p_m^i(k), \forall m \in \mathcal{N}(i)\right\}$. Condition (i) guarantees that if the gain of node i is at least as high as the perception of any of its neighboring nodes, then no agent starts moving away from node i at time k. This condition is required to guarantee the invariance of the desired distribution. Condition (ii) guarantees that there is at least ε_p agents at any node at any point in time, and is required so that conditions (iii) and (iv) are always well defined. Condition (iii) prevents there being too many agents that start moving away from node i at time k, so that the gain of node i just about reaches the highest perception of all of its neighboring nodes (of course additional agents could move to the neighboring node with the highest gain, reducing its value far enough, so that node i does actually become the node with the highest again at time $k+1$). Condition (iv) implies that if the gain of node i is less than the perception of some neighboring node, then at least a certain amount of agents (if not all but ε_p) must move to the neighboring node perceived as having the highest gain. Without condition (iv) some node with a high gain could be ignored by the agents and the desired distribution might never be achievable. Finally, note that satisfying conditions $(i)-(iv)$ requires that agents at node i know a_i, a_{j^*}, and b_{j^*} for some neighboring node $j^* \in \{j: \ p^i_j(k) \geq p^i_m(k), \ \ \forall m \in \mathcal{N}(i)\}$. We use $\gamma_{ij} \in \left(0, \frac{b_j}{a_i}\right) \subseteq (0,1)$ to characterize the degree of cooperation between agents at node j and those at node i. For agents at node j, the higher the value of γ_{ij}, the more willing they are to receive agents from node i, although this movement will degrade the gain of node j (i.e., if other agents do not leave node j). In Section IV we will see how these constants may be defined a priori in a specific application, in particular for a surveillance mission. In fact, their value will depend only on characteristics which are inherent to the group of agents and the environment being considered (e.g., the size of the region agents must cover and their moving capabilities).

(b) For all $e^{i \leftarrow \mathcal{N}(i)}_{\beta(i,k)} \in e(k)$, where $\beta(i,k) = (\beta_j(i,k): j \in \mathcal{N}(i))$ it is the case that:

$$0 \leq \beta_j(i,k) \leq \sum_{k'=k-B_t+1}^{k-1} \alpha_i(j,k') - \sum_{k'=k-B_t+1}^{k-1} \beta_j(i,k')$$

(c) If $e^{i \leftarrow \mathcal{N}(i)}_{\beta(i,k)} \in e(k)$ with $\beta_j(i,k) > 0$ for some j such that $j \in \mathcal{N}(i)$, then $e^{i \rightarrow \mathcal{N}(i)}_{\alpha(i,k)} \in e(k)$ with $\alpha(i,k) = (0,\ldots,0)$. In other words, if some agents arrive at node i at time k, then no agents will start moving away from that node i at the same time instant. Note that this assumption is not unrealistic in these types of problems, especially since it is imposed locally only.

- If $e(k) \in g(x(k))$, and $e^{i \rightarrow \mathcal{N}(i)}_{\alpha(i)}, e^{i \leftarrow \mathcal{N}(i)}_{\beta(i)} \in e(k)$, then $x(k+1) = f_{e(k)}(x(k))$, where

$$\begin{aligned} x_i(k+1) &= x_i(k) - \sum_{\left\{m:\, m \in \mathcal{N}(i),\, e^{i \rightarrow \mathcal{N}(i)}_{\alpha(i,k)} \in e(k)\right\}} \alpha_m(i,k) + \sum_{\{m:\, m \in \mathcal{N}(i)\,,\, e^{i \leftarrow \mathcal{N}(i)}_{\beta_m(i,k)} \in e(k)\}} \beta_m(i,k) \qquad (3) \\ x_{i \rightarrow j}(k+1) &= x_{i \rightarrow j}(k) + \alpha_j(i,k) - \beta_i(j,k) \end{aligned}$$

In other words, the amount of agents at node i at time $k+1$, $x_i(k+1)$, is the amount of agents at node i at time k, minus the total amount of agents leaving node i at time k, plus the total amount of agents reaching node i at time k. Note that Eq. (3) implies conservation of the amount of agents so that if $\sum_{i=1}^N x_i(0) = P$, $\sum_{i=1}^N x_i(k) = P$ for all $k \geq 0$. So if $x(0) \in \mathcal{X}$, $x(k) \in \mathcal{X}$, for all $k \geq 0$ (i.e., $\mathcal{X}$ is invariant).

Finally, we define a *partial event* to represent that *some* amount of agents start moving from $i \in \mathcal{H}$ to neighboring nodes $\mathcal{N}(i)$ and we denoted it by $e^{i \rightarrow \mathcal{N}(i)}$.

- For every substring $e(k), e(k+1), \ldots, e(k+B_s-1)$, there is the occurrence of partial event $e^{i \rightarrow \mathcal{N}(i)}$ for all $i \in \mathcal{H}$ (i.e., at any fixed time index k for all $i \in \mathcal{H}$ partial event $e^{i \rightarrow \mathcal{N}(i)} \in e(k')$ for some k', $k \leq k' \leq k + B_s - 1$). This restriction guarantees that by time $k + B_s - 1$ agents at any node $i \in \mathcal{H}$ must sense and potentially start moving according to conditions $(a)(i)-(iv)$.
- For every $i \in \mathcal{H}$, $j \in \mathcal{N}(i)$, and k' such that $e^{i \rightarrow \mathcal{N}(i)}_{\alpha(i,k')} \in e(k')$ and $\alpha_j(i,k') > 0$, there is some k'', $k' \leq k'' \leq k' + B_t - 1$ such that $e^{j \leftarrow N(j)}_{\beta(j,k'')} \in e(k'')$ and $\alpha_j(i,k') = \beta_i(j,k'')$. This restriction guarantees that agents that start moving at time k' arrive at their destination node by time $k' + B_t - 1$.

Let E_k denote the sequence of events $e(0), e(1), \ldots, e(k-1)$, and let the value of the function $X(x(0), E_k, k)$ denote the state reached at time k from the initial state $x(0)$ by application of the event sequence E_k. We now study the evolution of any $X(x(0), E_k, k)$ that is reached from any event sequence E_k where $e(0), e(1), \ldots, e(k-1)$ satisfy conditions $a(i) - (iv)$.

III. Results

The results in this section take into account that truncated nodes may emerge while agents are trying to achieve the desired distribution, i.e., the desired state of some nodes may equal ε_p. In particular, we consider the emergence of truncated nodes in graphs $\mathcal{G}(H, A)$ with a fully connected topology. We show that the desired distribution is an invariant set, and study its stability properties. Moreover, if we assume that the graph $\mathcal{G}(H, A)$ is no longer fully connected, we then show that for a large enough total number of agents P there are no truncated nodes at the desired distribution, but the same stability properties still hold.

Theorem 1: For a fully connected graph $\mathcal{G}(\mathcal{H}, \mathcal{A})$, unknown but bounded sensing and travel delays, and agents that satisfy conditions $(a)-(c)$, the invariant set $\mathcal{X}_b$ is exponentially stable. Moreover, there exists a constant $P_c > N\varepsilon_p$ such that if the total amount of agents is at least $P > P_c$, the invariant set $\mathcal{X}_b$ is exponentially stable for any connected graph $\mathcal{G}(\mathcal{H}, \mathcal{A})$.

(Due to space constraints we do not include the proof of Theorem 1 here. For detailed information about the proof the reader should contact the authors.) The authors in [12] study different load balancing problems under different types of load: discrete, continuous, and virtual. The analysis here considers the virtual load case, where the load (the agents) affects the different processors (the nodes) to different extents. By considering sensing and travel delays our analysis does not require that the real time between events e_k and e_{k+1} necessarily be greater than the greatest sensing plus the greatest travel delay. In this sense, unlike the virtual load model introduced in [12], we allow for a reduction of the degree of synchronicity forced upon the system. Furthermore, the results in Theorem 1 extend the virtual load case in [12] to allow for a non-linear mapping between the state and the virtual load, something that in the precence of delays has not been achieved before.

Exponential stability of the invariant set means that all agents are guaranteed to converge to $\mathcal{X}_b$ at a certain rate. The proof of Theorem 1 shows that in the *worst case* scenario the higher the value of $\gamma = \min_{ij}\{\gamma_{ij}\}$, the faster agents achieve convergence. Next, in Section IV we study the *average* behavior of the agents, and show how our model finds application in cooperative control problems.

IV. Application

Assume that the region of interest can be divided into equal-size areas, and our goal is to make the proportion of agents match the relative importance of monitoring each area. Assume that each node in $\mathcal{G}(\mathcal{H}, \mathcal{A})$ represents an area. Agents may travel from one area to another, but they will require up to $B_t - 1$ time steps to do so. Let us also assume that the number of areas is $N = 5$, $\varepsilon_p = 0$, $P = 100$, and $\mathcal{G}(\mathcal{H}, \mathcal{A})$ has a fully connected topology. Furthermore, for all $i \in \mathcal{H}$ we use gain functions of the form

$$s_i(x_i(k)) = R_i - r x_i(k)$$

where $x_i(k)$ is the amount of agents monitoring area i at time k (i.e., the amount of agents approaching or attending targets in that area), R_i is the average rate at which targets appear in area i, and r is the average rate of targets that an individual agent $\varepsilon_x \leq \varepsilon_p$ can attend (i.e., we assume that P can be expressed as $P = n\varepsilon_x$, where n is an arbitrarily large number which represents the total *number* of agents of size $\varepsilon_x > 0$). The size of an agent ε_x is arbitrarily small and is only defined to approximate the concept of an individual agent for the continuous model. The value of r depends only on the size of the areas and the maneuvering capabilities of any agent ε_x within an area. We assume that R_i determines the importance of monitoring area i, and so the higher its value, the more agents should be allocated there. The agents allocated in a particular area do not, however, know the values of $R_1, \ldots, R_5$ a priori, and can only sense or compute outdated gains of all other areas (based on information they receive or their own onboard sensors). Here, the gain level of an area represents the average rate at which targets have appeared in an area, but are not being or have not been attended by any agent. Note that the gains decrease linearly with the amount of agents. This relation results from assuming that agents monitoring the *same* area randomly approach any target located within that area (for a detailed discussion on different gain functions that may be used for surveillance see [15]).

Furthermore, note that since all areas have the same size and all agents the same maneuvering capabilities, $s_i((x_i))$ satisfies Eq. (1) with $a_i = b_i = r$ for all $i \in \mathcal{H}$. Thus, agents must only know that $a_i = a_{j^*} = b_{j^*} = r$ to verify conditions $a(i) - (iv)$. However, if they do not know the precise rate at which an individual agent can attend targets within the same area, they must define positive constants $\overline{a} > 0$ and $\underline{b} > 0$, such that $\overline{a} \geq \max\{a_i, a_{j^*}\}$ and $\underline{b} \leq b_{j^*}$, so that Eq. (1) still holds. While using $\gamma_{ij} = \underline{b}/\overline{a} \leq b_j/a_i < 1$ for all $i \in \mathcal{H},\ j \in \mathcal{N}(i)$, limits the maximum degree of cooperation between agents, our results in Section III guarantee that the agents will achieve the desired distribution as long as $\gamma_{ij} \subseteq (0, b_j/a_i)$. For our simulations we assume that agents know that the average rate at which an individual agent can attend targets is at least 0.08, so that $\underline{b} = 0.08$, and at most 1, so that $\overline{a} = 1$. We let $\gamma_{ij} = \gamma_{pq} = \gamma \leq 0.08$ for all areas.

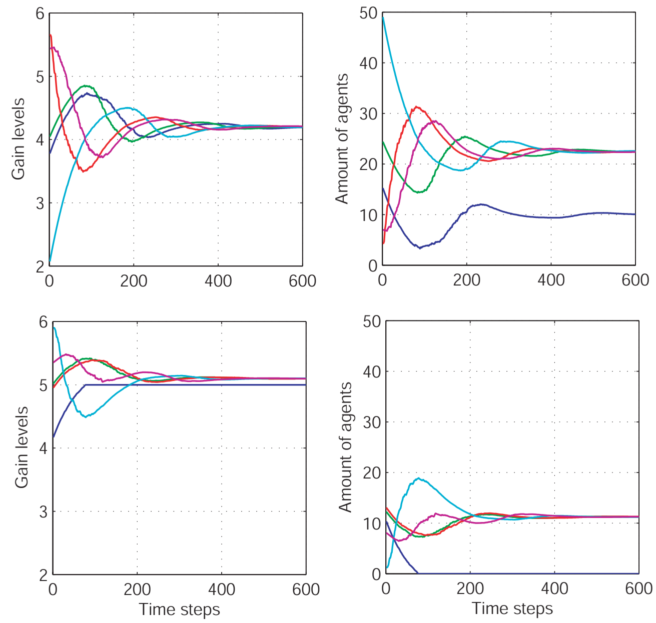

Fig. 1. Achieving the desired distribution for $N = 5$, $\gamma = 0.04$, $\varepsilon_p = 0$, $P = 100$ (top plots) and $P = 45$ (bottom plots) with random bounded sensing and travel delays.

Figure 1 shows the results of a sample run with $R_1 = \ldots = R_4 = 6$, and $R_5 = 5$, for 100 agents under random sensing and travel delays (bounded by $B_t = B_s = 10$). The top plots show that there are no truncated areas at the desired distribution. While all nodes achieve the same common gain, only about 10 agents are required in area 5 because it has a lower target appearance rate. Furthermore, if we let $P = 45$, the bottom plots show that the agents decide to cover only the areas with higher target appearance rates, and $s_5(0)$ remains below the achieved common gain.

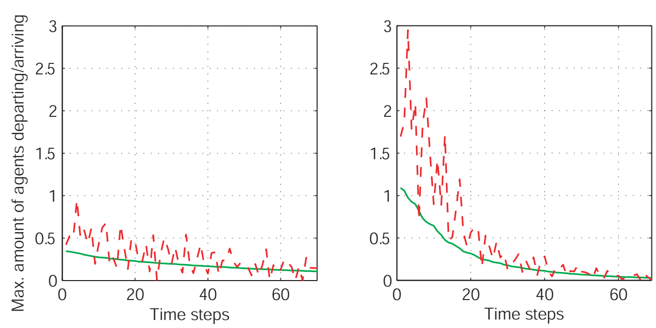

Fig. 2. Maximum amount of agents departing (solid curve) or reaching (dotted curve) any area when $\gamma = 0.02$ (left plot) and $\gamma = 0.04$ (right plot) with $B_t = B_s = 1$.

Figure 2 shows how the choice of γ affects the amount of agents traveling between areas. The left plot illustrates the maximum amount of agents that depart (solid curve) or reach (dotted curve) any area in the first 70 time steps when $\gamma = 0.02$. Since both maxima are key variables in determining the time needed to achieve the desired distribution, they must obviously decrease over time. Note that if we let $\gamma = 0.04$ agents behave more aggressively, in the sense that those in areas with higher gains are more willing to receive other agents from neighboring nodes, which leads to higher maxima but a faster convergence (see the right plot). In other words, Figure 2 seems to suggest that if speed of convergence to the desired distribution is more important than the cost of traveling (e.g., fuel expenditure), agents should cooperate as much as possible (i.e., γ should be made as large as possible).

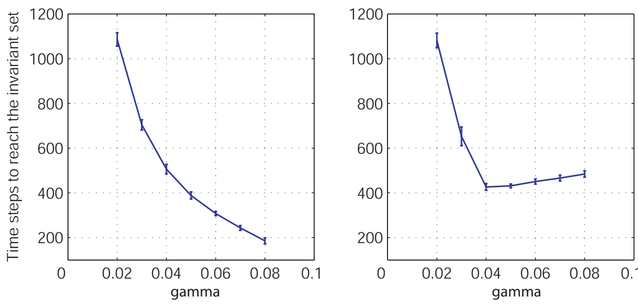

Fig. 3. The left plot shows the average time to reach $\mathcal{X}_b$ under no sensing or travel delays ($B_t = 2, B_s = 1$). The right plot illustrates the case when $B_t = 2$, but there are random sensing delays bounded by $B_s = 10$. Every data point represents 40 simulation runs with varying initial agent distributions, and the error bars are sample standard deviations for these runs.

We study this hypothesis in Figure 3 where we show the time required to achieve the desired distribution for different values of γ. The left plot shows the results under no sensing and travel delays (i.e., $B_t = 2, B_s = 1$), in which case higher values of γ lead to a faster convergence, as is also suggested from the worst case analysis (in the proof of Theorem 1). However, note that this relation no longer holds when sensing or travel delays are considered. The right plot in Figure 3 shows that the optimal cooperation level reduces from 0.08 to 0.04 when random sensing delays (bounded by $B_s = 10$) are considered, suggesting that a less aggressive behavior becomes desirable as the quality of the available information degrades. Thus, a higher degree of cooperation does not necessarily result, *on average*, in a faster convergence to the desired state.

Finally, the left plot in Figure 4 shows the average time when the desired distribution is reached for optimal values of γ and varying bounds on the sensing delays (i.e., for values of γ that *minimize* the average time to reach $\mathcal{X}_b$). It corroborates that the optimal degree of cooperation between agents depends on the quality of the information being used (i.e., the larger B_s, the less agents should cooperate in order to achieve the distribution the fastest). A similar plot results from considering travel delays, suggesting likewise that the longer it takes for agents to travel between different areas, the less they should cooperate. When comparing both types of delays, sensing delays seem to have a slightly larger effect

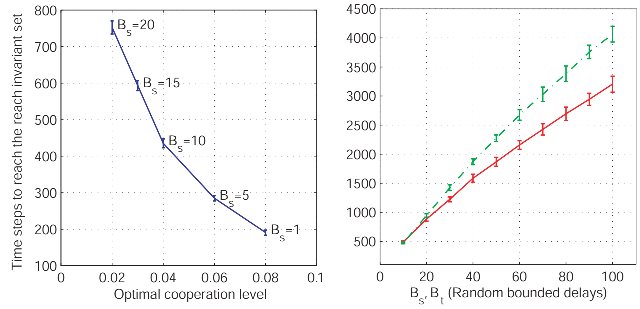

Fig. 4. The left plot shows the minimum (average) time to reach $\mathcal{X}_b$. The right plot shows the average time to reach $\mathcal{X}_b$ when vary B_s from 1 to 100 while keeping $B_t = 2$ (dotted curve), and vice versa (solid curve). For the right plot $\gamma = 0.02$. Every data point represents 40 simulation runs with varying initial agent distributions, and the error bars are sample standard deviations for these runs.

on achieving the desired distribution than travel delays (see the right plot in Figure 4 where we vary B_s from 1 to 100 while keeping $B_t = 1$ (dotted curve), and vice versa (solid curve)).

REFERENCES

[1] D. Bertsekas and J. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Belmont, Massachusetts: Athena Scientific, MA, 1997.

[2] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Autom. Control*, vol. 49, pp. 1520–1533, September 2004.

[3] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Trans. on Automatic Control*, vol. 50, no. 2, 2005.

[4] M. Cao, A. S. Morse, and B. D. Anderson, "Agreeing asynchronously: Announcement of results," in *Proceedings of the IEEE Conference on Decision and Control*, (San Diego, CA), December 2006.

[5] L. Wang and F. Xiao, "A new approach to consensus problems for discrete-time a new approach to consensus problems for discrete-time multiagent systems with time-delays," in *Proceedings of the 2006 American Control Conference*, (Minneapolis, MN), pp. 2118–2123, June 2006.

[6] Q. Hui and W. Haddad, "Distributed nonlinear control algorithms for network consensus," *Automatica*, vol. 44, p. 23752381, 2008.

[7] L. Fang, P. J. Antsaklis, and A. Tzimas, "Asynchronous consensus protocols: Preliminary results, simulations and open questions," in *Proceedings of the IEEE Conference on Decision and Control*, 2006.

[8] W. Ren, R. W. Beard, and E. M. Atkins, "A survey of consensus problems in multi-agent coordination," in *Proceedings of the American Control Conference*, (Portland, OR), pp. 3566–3571, June 2005.

[9] D. Angeli and P.-A. Bliman, "Stability of leaderless discrete-time multi-agent systems," *Mathematics of Control, Signals and Systems*, vol. 18, no. 4, pp. 293–322, 2006.

[10] A. Cortés, A. Ripoll, F. Cedó, M. A. Senar, and E. Luque, "An asynchronous and iterative load balancing algorithm for discrete load model," *J. Parallel Distrib. Comput.*, vol. 62, pp. 1729–1746, 2002.

[11] R. Elsässer and B. Monien, "Diffusion load balancing in static and dynamic networks," in *International Workshop on Ambient Intelligence Computing*, pp. 49–62, 2003.

[12] K. M. Passino and K. L. Burgess, *Stability Analysis of Discrete Event Systems*. John Wiley and Sons, Inc., NY, 1998.

[13] K. L. Burgess and K. M. Passino, "Stability analysis of load balancing systems," *Int. Journal of Control*, vol. 61, pp. 357–393, February 1995.

[14] J. Finke and K. M. Passino, "Stable cooperative multiagent spatial distributions," in *Proceedings of the IEEE Conference on Decision and Control and the European Control Conference*, (Seville, Spain), pp. 3566–3571, December 2005.

[15] J. Finke and K. M. Passino, "Stable cooperative vehicle distributions for surveillance," *ASME Journal of Dynamic Systems, Measurement, and Control*, vol. 129, pp. 597–608, September 2007.