Proceedings of the
44th IEEE Conference on Decision and Control, and
the European Control Conference 2005
Seville, Spain, December 12-15, 2005

Stable Cooperative Multiagent Spatial Distributions

Jorge Finke and Kevin M. Passino
Department of Electrical and Computer Engineering
The Ohio State University
finkej@ece.osu.edu, passino@ece.osu.edu

Abstract— **This paper introduces a mathematical model of the behavior of a group of agents and their interactions in a shared environment. We represent environmental spatial constraints that allow us to model range-limited sensing, motion, and communication capabilities of the agents. We derive general sensing, coordination, and motion conditions on the agents that guarantee that an "ideal free distribution" (IFD) of the group of agents will emerge across the environment. We show the impact of group size on the distribution of agents, and consider the emergent distribution for different classes of environments. Finally, we show how this theory is useful in solving a multivehicle cooperative surveillance problem.**

I. Introduction

The ideal free distribution concept from ecology characterizes how animals optimally distribute themselves across habitats [1], [2]. The word "ideal" refers to the assumption that animals have perfect sensing capabilities for simultaneously determining habitat "suitability" (assumed to be a correlate of Darwinian fitness) in each of a finite number of habitats. The ideal assumption also supposes that each animal will move to maximize its fitness (i.e., it moves to the habitat that is best for it). "Free" indicates that animals can move at no cost and instantaneously from any habitat directly to any other habitat at any time. The IFD pattern results from perfectly informed unconstrained local decisions by multiple animals. If an animal perceives one habitat as more suitable, via some correlate of fitness such as the rate of arrival of nutrients, it can move to it. This movement will, however, reduce the new habitat's suitability, both to itself and other animals in that habitat. The IFD is an equilibrium distribution where no animal can increase its fitness by unilateral deviation from one habitat to another; hence at the IFD all animals achieve equal fitness and the IFD is a Nash equilibrium.

Many extensions of the IFD have been developed [3]. An important extension to the model takes into account that individuals differ in competitive ability, as in [4],[5]. Other work that focuses on competitiveness can be found in [6],[7]. The author in [7] introduces the concept of interference as the direct effect caused by the presence of several competitors in the same habitat. The IFD model we introduce here is different from any existing ones in the literature. It is built on a directed graph. We use a generic terminology for IFD concepts around this graph, one that is appropriate for biology and engineering. We refer to habitats, food sources, resource sites, etc. as nodes. Each node is characterized by its suitability, which represents how profitable or suitable the node appears to an agent. It could be that suitability represents a task input rate to a node, where there is a certain value to having an agent process each of these tasks. For an animal, suitability could represent the consumption rate achieved at a habitat, or the probability of finding a mate or shelter there. The graph topology defines the graph's interconnections between nodes via a set of directed arcs. Hence, the graph topology allows us to represent removal of both the ideal and free restrictions to the original IFD model. We do require, however, minimal restrictions on the graph topology to ensure that an IFD can be achieved.

Our model focuses on the individual agents' motion dynamics across the graph that drive the behavior of the group as a whole. We consider a general class of habitat suitability functions. We show how an "invariant set" can represent the IFD. Theorems 1 and 2 give properties of this invariant set in terms of the agent group size and connectedness of the habitats. We then use Lyapunov stability analysis of the invariant set to illustrate that there is a wide class of agent strategies (i.e., "proximate" decision-making mechanisms), and resulting agent movement trajectories across nodes, that still achieve the desired distribution. In particular, Theorem 3 shows that the IFD is asymptotically stable in the large for any graph topology with a large enough agent population size. Theorem 4 shows that the IFD is asymptotically stable in the large for any population size and a fully connected graph topology (e.g., an environment where agents can sense and move to any part of the environment regardless of their current location). Next, in Theorem 5 we show that if we constrain the level of asynchronicity in agent decision-making, the IFD is exponentially stable in the large and the rate at which the IFD is achieved can be quantified in terms of characteristics of the connectedness of the habitats. The results extend the existing theory of the IFD by showing the impact of a class of suitability functions, agent perceptual constraints, travel constraints, movement trajectories, and agent decision-making strategies on achievement of the IFD. Our analysis shows how a global distribution pattern can *emerge* from poorly informed and constrained individual decision-making, something that is of significant interest in biology (e.g., see [4]).

This work was supported by the AFRL/VA and AFOSR Collaborative Center of Control Science (Grant F33615-01-2-3154). We gratefully acknowledge the help from Andy Sparks and Corey Schumacher at AFRL in formulating the cooperative surveillance problem. We also gratefully acknowledge the inputs of Thomas Waite, Alvaro Gil, Brandon Moore and Nicanor Quijano. Please address correspondence to K. M. Passino, 416 Dreese Laboratory, 2015 Neil Avenue, Columbus, Ohio 43210-1272.

0-7803-9568-9/05/$20.00 ©2005 IEEE

In cooperative control, there is a significant amount of current research actively focused on how to allocate vehicles to tasks in order to minimize a performance metric. Such work includes the study of distributed load balancing [8], "persistent area denial (PAD)" [9], [10], and applications of the m-person Dynamic Traveling Repairperson Problem (m-DTRP) to cooperative control [11]. Here, we show how to use the IFD concept and our theoretical extensions to solve a type of cooperative surveillance problem. We show how an individual IFD-based decision-making strategy is useful for coordinating the distribution of multiple vehicles over an environment with random pop up targets. In particular, the environment is divided into different regions and vehicles distribute themselves to achieve a similar target service rate in all regions.

It is, however, the general idea behind strategies for IFD achievement that is most important to the area of cooperative control. In most of the studies cited above (e.g., [9], [10], [11] and many others) there is a coupling between assignment of spatially distributed tasks and spatial position of all vehicles. When a vehicle is assigned an area to perform tasks, the benefit of assigning this area to all other vehicles decreases. The IFD strategies show a way to untangle this space and time dependent "multi-vehicle to multi-area (-task)" coupling to provide good cooperative behavior.

II. Agents and Environment

A. Habitat Suitability

Assume that there are N habitats (nodes). Define the suitability of node i as $s_i(x_i)$, where $x_i \geq \varepsilon_p$ is a scalar that represents the amount of agents at node i, $i = 1, 2, \ldots, N$ (converting x_i to the density of agents at habitat i is achieved via linearly scaling x_i), and $\varepsilon_p \geq 0$ is the minimum amount of agents allowed at any node. Assume the following:

- *Fixed agent population:* Let $\sum_{i=1}^{N} x_i = P$, where $P > N\varepsilon_p$ is a constant so there are a fixed amount of agents in the environment.
- *Suitability decreases with an increasing amount of agents:* Assume that the s_i, $i = 1, \ldots, N$, satisfy

$$\frac{s_i(y_i) - s_i(z_i)}{y_i - z_i} \leq -c < 0 \tag{1}$$

  for any $y_i, z_i \in [\varepsilon_p, P]$, $y_i \neq z_i$. Thus, $s_i(x_i)$ is a strictly monotone decreasing function in x_i. We also assume that $\lim_{x_i \to \infty} s_i(x_i) = 0$ for all $i = 1, \ldots, N$.
- *Suitability changes are related to changes in the amount of agents:* We assume that all suitability functions $s_i(x_i)$, $i = 1, \ldots, N$, satisfy a Lipschitz condition on $[\varepsilon_p, P]$; that is, for every node i there exists a constant $K_i > 0$ such that

$$\frac{|s_i(y_i) - s_i(z_i)|}{|y_i - z_i|} \leq K_i \tag{2}$$

  for any $y_i, z_i \in [\varepsilon_p, P]$, $y_i \neq z_i$. This eliminates the possibility that a very small change in the amount of agents at a habitat can result in a large change in suitability.
- *Strictly positive suitability:* We assume that the functions $s_i(x_i) > 0$ for all $i = 1, \ldots, N$, and all $x_i \in [\varepsilon_p, P]$.

B. Environmental Constraints on Agent Sensing and Motion

We will consider a general graph topology to model interconnections between nodes. The nodes are $H = \{1, ..., N\}$. The interconnection of nodes is described by a directed graph, (H, A), where $A \subset H \times H$. If $(i, j) \in A$, then this represents that an agent at node i can sense node j and can move from i to j. For agents at node i, where $(i, j) \in A$, "sensing node j" implies that agents at node i know $s_j(x_j)$ and x_j. If the function s_j were known by agents at node i and invertible it would be sufficient to sense the value of $s_j(x_j)$ only. For every $i \in H$, there must exist some $j \in H$, $i \neq j$, such that $(i, j) \in A$ and there exists a path between any two nodes, in order to ensure that every node is connected to the graph. We also assume that if $(i, j) \in A$ then $(j, i) \in A$ so that if an agent is at i and can move to j (sense the suitability at j), agents at j can also move from j to i (sense the suitability at i, respectively). An agent at node i can only directly move to node j if $(i, j) \in A$. However, if $(i, j) \notin A$ it may in some situations be possible for an agent to (indirectly) move to j by passing through a series of other nodes. If $(i, j) \in A$, then $i \neq j$; however, agents at i know the value of $s_i(x_i)$, are assumed to know x_i, and are already at i so they do not need to move to get to it.

We use the distributed discrete event system modeling methodology from [12]. Let $\mathbb{R}_{\geq \varepsilon_p} = [\varepsilon_p, \infty)$ and $\Delta = \left\{ x \in \mathbb{R}^N_{\geq \varepsilon_p} : \sum_{i=1}^{N} x_i = P \right\}$ be the simplex over which the x_i dynamics evolve. Constraints on our model below will ensure that $x(k) \in \Delta$ for all $k \geq 0$. Let $\mathcal{X} = \Delta$ be the set of states. Let $x(k) = [x_1(k), x_2(k), ..., x_N(k)]^\top \in \mathcal{X}$ be the state vector, with $x_i(k)$ the amount of agents at node i at time index $k \geq 0$. Here we assume that there is a high number of agents so that x_i is accurately represented as a continuous variable, a common approach in theoretical ecology. Let $I(x) = \{i \in H : x_i > \varepsilon_p, \ x \in \Delta\}$ represent the set of nodes that are occupied (inhabited) by more than ε_p agents, and let $U(x) = H - I(x)$ represent the set of nodes that are uninhabited for state x. The size of the set $I(x)$ is denoted by N_I. Figure 1 shows an example of a system with $N = 3$ nodes. Note that any horizontal line crossing at least one s_i curve represents an IFD state for some P.

C. Agent Sensing, Coordination, and Motion Requirements

To illustrate individual agent sensing, coordination, and motion requirements we assume that the agent population size 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" $\varepsilon_x > 0$. We do not assume that an agent has a fixed size ε_x. The size of an agent ε_x is arbitrarily small and is only defined to approximate the concept of an individual agent for a continuous model. This allows us to discuss coordination at two different levels: by coordination at the "agent level" we mean information sharing, agent conditions, and agreement strategies between agents that are

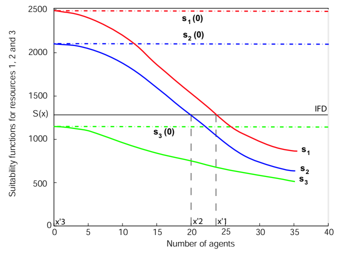

Fig. 1. Suitability functions $s_i(x_i)$ for three fully connected nodes. The IFD distribution is reached when all agents have exactly the same suitability in the habitats they are in. After the IFD is reached there is no movement of agents between nodes. For the example shown in the plot, while agents distribute themselves over nodes 1 and 2, node 3 remains without agents since its suitability s_3 is too low to be chosen by any agent.

located at the *same* node (e.g., between those individuals sharing the same habitat). By coordination at the "node level" we mean node conditions that may be satisfied with respect to neighboring nodes. In general, coordination at the node level is achieved via coordination at the agent level (unless there was some physical mechanism to provide for node-level coordination). While we assume that every agent can communicate, share information, and coordinate with every other agent within the same node (at the agent level), coordination strategies at the node level depend on the topology of the graph (H, A). We now specify the discrete event model and define agent sensing, coordination, and motion requirements at this level.

Let $\mathcal{E}$ be a set of events and let $e^{i,p(i)}_{\alpha(i)}$ represent the event that one or more agents move from node $i \in H$ to *neighboring* nodes $m \in p(i)$, where $p(i) = \{j : (i,j) \in A\}$. Movement of agents from node i to neighboring nodes decreases x_i. Let $\alpha_m(i)$ denote the quantity of agents that move from node $i \in H$ to node $m \in p(i)$. Let the list $\alpha(i) = (\alpha_j(i), \alpha_{j'}(i), \ldots, \alpha_{j''}(i))$ such that $j < j' < \cdots < j''$ and $j, j', \ldots, j'' \in p(i)$ and $\alpha_j \geq 0$ for all $j \in p(i)$; the size of the list $\alpha(i)$ is $|p(i)|$. For convenience, we will denote this list by $\alpha(i) = (\alpha_j(i) : j \in p(i))$. Let $\{e^{i,p(i)}_{\alpha(i)}\}$ denote the set of *all* possible combinations of how agents can move between nodes (i.e., $\alpha(i) \in \mathbb{R}^{|p(i)|}_{\leq P}$, where $\mathbb{R}_{\leq P} = [0, P]$). Let the set of events be described by $\mathcal{E} = \mathcal{P}(\{e^{i,p(i)}_{\alpha(i)}\}) - \{\emptyset\}$ ($\mathcal{P}(\cdot)$ denotes the power set). Notice that each event $e(k) \in \mathcal{E}$ is defined as a *set*, with each element of $e(k)$ representing the transition of possibly multiple agents among neighboring nodes in the graph. Multiple elements in $e(k)$ represent the simultaneous movements of agents, i.e., migrations out of multiple nodes.

An event $e(k)$ may only occur if it is in the set defined by an "enable function," $g : \mathcal{X} \longrightarrow \mathcal{P}(\mathcal{E}) - \{\emptyset\}$. State transitions are defined by the operators $f_e : \mathcal{X} \longrightarrow \mathcal{X}$, where $e \in \mathcal{E}$. Let $\gamma_{ij} \in (0,1)$ for $(i,j) \in A$ represent the proportion of imbalance in nodes' suitability that is *sometimes* guaranteed to be reduced when agents move from node i to node j. We now specify g and f_e for $e(k) \in g(x(k))$, which define the agents' sensing and motion:

- If for a node $i \in H$, $s_i(x_i) \geq s_j(x_j)$ for all $(i,j) \in A$, then $e^{i,p(i)}_{\alpha(i)} \in e(k)$ such that $\alpha(i) = (0, \ldots, 0)$ is the only enabled event. Hence, agents at the most suitable node that they know of do not move. Note also that this does not then allow for a "swap" of equal number of agents between two nodes i and j, $(i,j) \in A$ such that $s_i(x_i) = s_j(x_j)$.
- If for node $i \in H$, $s_i(x_i) < s_j(x_j)$ for some j such that $(i,j) \in A$, then the only $e^{i,p(i)}_{\alpha(i)} \in e(k)$, are ones with $\alpha(i) = (\alpha_j(i) : j \in p(i))$, such that:

$$
\begin{aligned}
&(i) && x_i - \sum_{m \in p(i)} \alpha_m(i) \geq \varepsilon_p, \\
&(ii) && s_i\left(x_i - \sum_{m \in p(i)} \alpha_m(i)\right) \leq s_{j^*}\left(x_{j^*} + \alpha_{j^*}(i)\right) \\
& && \text{for some } j^* \in \{j : \ s_j(x_j) \geq s_m(x_m), \\
& && \text{for all } \ m \in p(i)\}, \ \text{ and} \\
&(iii) && s_{j^*}\left(x_{j^*} + \alpha_{j^*}(i)\right) \\
& && \leq s_{j^*}(x_{j^*}) - \gamma_{ij^*}(s_{j^*}(x_{j^*}) - s_i(x_i)) \\
& && \text{for some } j^* \in \{j : \ s_j(x_j) \geq s_m(x_m), \\
& && \text{for all } \ m \in p(i)\}
\end{aligned}
$$

  Condition (i) guarantees that at any node there are at least ε_p agents. It is required so that conditions (ii) and (iii) are well defined at all times. Conditions (ii) and (iii) constrain how agents can move in terms of node suitabilities. Condition (ii) relaxes to a certain extent the "ideal" part of the IFD assumptions, since node i becoming better than some of its neighbors means that some agents moved from node i to neighboring nodes that might have seemed promising, but resulted in lower suitabilities than node i (e.g., because many agents might have simultaneously decided to leave node i). Condition (iii) also lifts to a certain extent the "ideal" part of the IFD assumptions, since it allows some agents to move to nodes that do not necessarily correspond to a best suitability choice, as long as at least some individuals do. Condition (ii) together with condition (iii) guarantees that the highest suitability node is strictly monotone decreasing over time.
- If $e(k) \in g(x(k))$, $e^{i,p(i)}_{\alpha(i)} \in e(k)$, then $x(k+1) = f_{e(k)}(x(k))$, so that $x_i(k+1)$ equals $x_i(k)$ plus

$$
\sum_{\{j:\, i \in p(j)\ ,\ e^{j,p(j)}_{\alpha(j)} \in e(k)\}} \alpha_i(j) - \sum_{\{j:\, j \in p(i),\ e^{i,p(i)}_{\alpha(i)} \in e(k)\}} \alpha_j(i)
$$

  Note that the definition of $f_{e(k)}(x(k))$ implies conservation of the number of agents so that if $x(0) \in \Delta$, $x(k) \in \Delta$, $k \geq 0$ (i.e., Δ is invariant).

Specifying g and f_e for $e(k) \in g(x(k))$ at the node level allows for a wide class of interagent coordination strategies at the agent level. Moreover, different nodes may even have different coordination strategies at the agent level. In general, if there is no coordination at the agent level at all, then the above conditions may not be satisfied. Deciding where to go is not a decision made by an agent based on its own assessment of the neighboring nodes only, but it must consider how the other agents at the same node behave. In other words, agents must know (e.g., communicate, coordinate with) how other agents at the same node plan to move. Note that communication between agents must not necessarily mean true signaling and information transfer between agents. It may also be based on "cues" in the environment (i.e., if an agent intending to leave a node observes another agents that intend to leave the node as well, it may wait to see how their migration affects the suitabilities of the neighboring nodes before deciding on where to go). Therefore, an agent's decision on where to go must be made relative to where other agents at the same node decide to go.

Let $\mathcal{E}^{\mathbb{N}}$ denote the set of all infinite sequences of events in $\mathcal{E}$. Let $E_v \subset \mathcal{E}^{\mathbb{N}}$ be the set of valid event trajectories for the model (i.e., ones that are physically possible). Event $e(k) \in g(x(k))$ is composed of a set of what we will call "partial events." Define a *partial event of type* i to represent the movement of $\alpha(i)$ agents from node $i \in H$ to its neighbors $p(i)$ so that conditions (i) - (iii) are satisfied. A partial event of type i will be denoted by $e^{i,p(i)}$ and the occurrence of $e^{i,p(i)}$ indicates that some agents located at node $i \in H$ move to other nodes. Partial events must occur according to the "allowed" event trajectories. The allowed event trajectories define the degree of asynchronicity of the model at the node level. We define two possibilities for the allowed event trajectories:

(i) For allowed event trajectories $E_i \subset E_v$, assume that each type of partial event occurs infinitely often on each event trajectory $E \in E_i$. The assumption is met if at each node all agents do not ever stop trying to move (e.g., if each agent persistently tries to move to neighboring nodes). This corresponds to assuming "total asynchronism" [13].

(ii) For allowed event trajectories $E_B \subset E_v$, assume that there exists $B > 0$, such that for every event trajectory $E \in E_B$, in every substring $e(k'), e(k'+1), e(k'+2), \ldots, e(k'+(B-1))$ of E there is the occurrence of every type of partial event (i.e., for every $i \in H$, the partial event $e^{i,p(i)} \in e(k)$, for some k, $k' \leq k \leq k'+B-1$). This corresponds to assuming "partial asynchronism" [13]. Loosely speaking, the assumption is met if agents try to move to neighboring nodes every certain number of steps. It is by no means assumed that the time index k is known to all agents. Instead, k should be viewed as a time index seen by an external observer holding a global clock. Each node and each agent can be viewed as obeying its own local clock that is not arbitrarily out of synch with the global clock. If we assume that these node and agent level clocks can neither be arbitrarily fast nor arbitrarily slow relative to the global clock and that partial events occur at least once during a time interval of finite length as measured by a node's clock, we can ensure the existence of B [13]. Coordination strategies at the agent level in node i must therefore guarantee the occurrence of a partial event of type i at least once during a time interval of finite length by its own clock to ensure that B will exist.

Finally, 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.

III. Emergent Agent Distributions

The set

$$\begin{aligned}\mathcal{X}_c \quad = \quad & \{\, x \in \mathcal{X} : \text{ for all } i \in H, \ \textit{either } s_i(x_i) = s_j(x_j) \\ & \text{for all } j \in p(i) \text{ such that } x_j \neq \varepsilon_p \text{ and } s_i(x_i) \geq s_j(x_j) \\ & \text{for all } j \in p(i) \text{ such that } x_j = \varepsilon_p, \textit{ or } x_i = \varepsilon_p \}\end{aligned}$$

represents a distribution of agents. Any distribution $x \in \mathcal{X}_c$ is such that for any $i \in H$ either $x_i = \varepsilon_p$, in which case node i has the minimum amount of agents allowed at that node; or if $x_i \neq \varepsilon_p$ it must be the case that all neighboring nodes $j \in p(i)$ such that $x_j \neq \varepsilon_p$ have the same suitability levels as node i. In $\mathcal{X}_c$, if $x_j = \varepsilon_p$, for $j \in p(i)$, then $s_j(x_j) \leq s_i(x_i)$. Notice that the only $e(k) \in g(x(k))$, when $x(k) \in \mathcal{X}_c$, are ones such that all $e^{i,p(i)}_{\alpha(i)} \in e(k)$ have $\alpha(i) = (0, 0, \ldots, 0)$ since conditions $(i) - (iii)$ cannot be satisfied. Hence, if $x(k) \in \mathcal{X}_c$, $x(k') = x(k)$ for $k' \geq k$ (i.e., when all motion conditions are satisfied, $\mathcal{X}_c$ is an invariant set under the flow of the system). Recall that for any $x \in \mathcal{X}_c$ there exists a set of inhabited nodes we denote by $I(x) \subseteq H$. The size or composition of $I(x)$ and the achieved suitability levels in $\mathcal{X}_c$ are *not* always known a priori or at any point before the set $\mathcal{X}_c$ is reached. The set $I(x)$ and the suitability levels emerge while agents distribute themselves over the nodes.

A. Emergence of Habitat Patches

Note that according to the definition of $\mathcal{X}_c$ it is possible for unconnected nodes (i.e., ones such that $(i, j) \notin A$) in the set $I(x)$ to have *different* suitabilities when the distribution is achieved. This could happen if two inhabited nodes with high suitabilities are separated by an uninhabited node. However, any two nodes that are linked according to the graph (H, A) (i.e., ones such that $(i, j) \in A$) *and* belong to the set $I(x)$ must have the same suitability at the desired distribution. Hence, depending on the graph's connectivity, there could be isolated "habitat patches" of inhabited nodes where only nodes belonging to the same habitat patch have equal suitability (i.e., forming an environment of different habitat patches). Moreover, note that the formation of habitat patches depends on the total number of agents, their initial distribution $x(0)$, and random events.

B. Agent Distribution Properties

Notice that in general there are many different agent distributions such that $x \in \mathcal{X}_c$. Indeed, for an arbitrary environment, the number of potential IFDs, $|\mathcal{X}_c|$, is in general infinite, even for a fixed number of agents. We now restrict our analysis to IFD distributions where one habitat patch may emerge only. We first show some properties of the invariant set that will be useful in the later analysis of the agents' dynamics in Section IV.

Theorem 1 (No Truncation Case): Given (H, A) there exists a constant $P > N\varepsilon_p$, such that if the total amount of agents is at least P, then the invariant set $\mathcal{X}_c$ satisfies $|\mathcal{X}_c| = 1$, and all nodes are inhabited at the IFD so that $I(x) = H$ for $x \in \mathcal{X}_c$.

(Due to space constraints we do not include any proofs here. For detailed information about the proofs of any of the theorems the reader should contact the authors.) Theorem 1 states that for a large enough group of agents, there are no truncated nodes at the IFD. In other words, $U(x) = \emptyset$ for all $x \in \mathcal{X}_c$. Hence, every node must have the same suitability at the IFD. Moreover, since $|\mathcal{X}_c| = 1$, the agent distribution at the IFD is unique. In other words, for any initial agent distribution there exists only one distribution that belongs to $\mathcal{X}_c$ and represents the IFD.

Next, let us assume that we have a fully connected graph topology (i.e., every node connects to every other node). This is consistent with the assumptions in [1], [2].

Theorem 2 (Truncation Case): For a fully connected network (H, A) and any population size P, the invariant set $\mathcal{X}_c$ satisfies $|\mathcal{X}_c| = 1$.

Theorem 2 implies that for a fully connected graph topology there is only one habitat patch of inhabited nodes independent of the amount of agents P. The full connectedness of the habitats leads to suitability equalization across all inhabited nodes and the emergence in some cases of a set of uninhabited nodes (e.g., if the population size P is not large enough). Given the assumptions of Theorems 1 and 2 it may be possible for some cases to explicitly find $x \in \mathcal{X}_c$. Our analysis below, however, is not dependent on knowing the explicit $x \in \mathcal{X}_c$.

IV. Stability Analysis: Emergent Distribution

Section III studied the characteristics of the invariant set that represents the IFD distribution for different population sizes and connectedness characteristics of environments. We now study how the group of agents approach this set.

A. Asymptotic Stability of the IFD

Let us consider again a general graph topology (H, A) and assume that every node is connected to the graph, but not every node connects to every other node.

Theorem 3 (Asymptotic Stability of the IFD): Given (H, A), $\varepsilon_p \geq 0$, and agent motion conditions $(i)-(iii)$, there exists a constant $P > N\varepsilon_p$ such that if the total amount of agents is at least P, then the invariant set $\mathcal{X}_c$ is asymptotically stable in the large with respect to E_i.

Since $\mathcal{X}_c$ is asymptotically stable *in the large*, there is only one equilibrium distribution for each population of at least P agents. Thus, for any initial agent distribution this equilibrium will be achieved. Note that this result provides general sufficient conditions on when an IFD is achieved. Moreover, our analysis considers all environments which can be modeled by a wide class of suitability functions. It includes functions which have been found to be useful in biology, like the one originally used to introduce the IFD concept in [1], and the one in [5] which introduced the interference model, among others. We also extend the existing IFD theory by considering a general interconnection topology, which allows us to consider less restrictive agent sensing and motion abilities. Theorem 3 is an extension of the load balancing [13] theorems in [14], [12] to the case when the "virtual load" is a nonlinear function of the state.

B. Emergence of Uninhabited Habitats

Let us now consider an unconstrained environment and an arbitrary number of agents.

Theorem 4 (Asymptotic Stability of the IFD, Emergence of Uninhabited Habitats): For a fully connected network (H, A), $\varepsilon_p > 0$, any population size P, and agent motion conditions $(i)-(iii)$, the invariant set $\mathcal{X}_c$ is asymptotically stable in the large with respect to E_i.

Notice that Theorem 4 requires $\varepsilon_p > 0$ because if $\varepsilon_p = 0$ at a truncated node i, then $s_i(x_i)$ equals infinity for certain suitability functions (e.g., $s_i(x_i) = \frac{a_i}{x_i}$). The proof of Theorem 4 considers the emergence of different habitat patches when the environment is modeled by a fully connected topology. Habitat patches emerge as agents distribute themselves over the nodes, and the total population size is small enough. The dynamic emergence of habitat patches is considered in the proof of Theorem 4, and is something that had not been analyzed in the literature before.

C. Rate of Convergence to the IFD

We now assume more restrictive sensing and motion conditions in order to study the rate of convergence to the desired distribution. In particular, let us assume that the allowed event trajectories are E_B so that agents at any node will try to move to neighboring nodes at least every B steps. We consider again a general graph topology (H, A) and assume that every node is connected to the graph, but not every node connects to every other node.

Theorem 5 (Exponential Stability of the IFD): Given (H, A), $\varepsilon_p \geq 0$, and agent motion conditions $(i)-(iii)$, there exists a constant $P > N\varepsilon_p$ such that if the total number of agents is at least P, then the invariant set $\mathcal{X}_c$ is exponentially stable in the large with respect to E_B.

Note that $\mathcal{X}_c$ is not exponentially stable with respect to E_i; the guarantee of occurrence of partial events with B is critical. Exponential stability of an invariant set means that all agents are guaranteed to converge to $\mathcal{X}_c$ at a certain rate. In particular, as B increases the rate of convergence to $\mathcal{X}_c$ decreases, since agents are only guaranteed to move at a slower migration rate. Furthermore, following [13], let $R = \max_i\{|p(i)|\}$ be the maximum number of neighboring nodes for any node $i \in H$. Then, if we assume a fully connected graph topology, the proof of Theorem 5 shows that the highest suitability of all nodes must decrease every RB steps. Note also that according to Equation (1), the constant $c > 0$ represents an intrinsic characteristic of the class of suitability functions being considered. In particular, a large value of c classifies a set of functions whose suitabilities decrease quickly with an increasing amount of agents. Note that a larger value of c guarantees, in general, a faster converge to $\mathcal{X}_c$.

V. APPLICATION: COOPERATIVE SURVEILLANCE

Here, we assume that there are $P = 16$ vehicles (with $\varepsilon_x = 1$), which must cover an environment defined by a square area of $100\ \mathrm{km}^2$ (hence x_i is not a continuous variable). We assume that the environment is divided into four equally sized regions. Every node $i \in H$ represents a region. Let us assume that every four seconds a new pop up target randomly appears anywhere in the environment. We also assume that vehicles can move from any region to any other region and consider therefore a fully connected graph (H, A). Moreover, vehicles have complete information about the environment and may even share information with vehicles that are not necessarily in the same region. In particular, we assume that the location of targets which have appeared in the environment is known to every vehicle (e.g., via satellite information). Suitability functions for every region are defined as the overall rate of appearance of unattended targets. In particular, the suitability of node i is defined as the number of targets present in that region that are not being or have not been visited by any vehicle in a time window divided by the length of that window (hence this is an approximation of the overall rate of appearance of unattended targets). If a vehicle approaches a target located in region i, this will decrease the suitability in that region and increase the overall target service rate. Our goal is to achieve similar overall target service rates in all regions.

In order to evaluate different vehicle strategies, we define the mission performance measure as the time needed for the difference between any two suitability levels to reach and settle within a given range (here 4%). We denote the settling time for a given mission by t_s. We will first compare the performance of the IFD-based agent strategy to a greedy strategy. The left plot in Figure 2 shows how suitability levels change over time during the first 200 seconds of a mission. It represents the case when vehicles violate the proposed conditions $(i) - (iii)$, and simply approach the region that seems to be the most suitable for them (e.g., the one with the highest rate of appearance of unattended targets). Note that the suitabilities levels do not converge to any particular any value. On the other hand, the right side plot in Figure 2 represents the case when vehicles distribute themselves over the environment while satisfying conditions $(i) - (iii)$. Here, we assume that $\gamma_{ij} \approx 0$ for all $(i, j) \in A$ so that condition (iii) can easily be satisfied by any single vehicle. Note that the suitability levels converge and the settling time in this case is approximately $t_s = 100$ seconds. Moreover, for any time $t \geq t_s$ the overall rate of appearance of unattended targets in all regions differs by less than 4%.

Next, we perform Monte Carlo runs to compare different vehicle strategies at the agent level. In particular, we want to show the effect of different cooperation levels between vehicles. Here, we assume that vehicles cooperate by sharing information about where to go. Figure 3 shows how the mean settling time decreases with the number of vehicles any vehicle may cooperate with. It is interesting to note that for lower numbers of cooperating vehicles the settling times

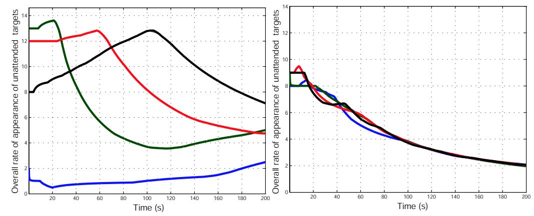

Fig. 2. Suitability levels; greedy strategy(left), cooperative strategy(right).

do not vary much; however, as the number of cooperating vehicles increases, the effect of adding an additional cooperating vehicle is more noticeable.

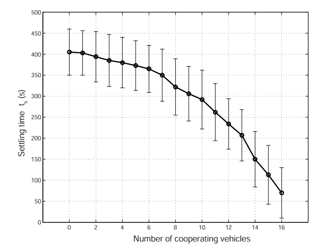

Fig. 3. Settling time for different coordination strategies at the agent level. Every data point represents 60 simulation runs with varying target pop up locations. The error bars are standard deviations for these runs.

REFERENCES

[1] S. D. Fretwell and H. L. Lucas, "On territorial behavior and other factors influencing distribution in birds," *Acta Biotheoretica*, vol. 19, pp. 16–36, 1970.

[2] G. A. Parker, "The reproductive behaviour and the nature of sexual selection in scatophaga stercoraria," *Evolution*, vol. 28, pp. 93–108, 1974.

[3] T. Tregenza, "Building on the ideal free distribution," *Advances in Ecological Research*, vol. 26, pp. 253–307, 1995.

[4] W. J. Sutherland, *From Individual Behaviour to Population Ecology*. New York: Oxford University Press, 1996.

[5] G. A. Parker and W. J. Sutherland, "Ideal free distributions when individuals differ in competitive ability: phenotype-limited ideal free models," *Animal Behavior*, vol. 34, pp. 1223–1242, 1986.

[6] M. Milinski, "Competitive resource sharing: An experimental test for a lerning rule for ESSs," *Animal Behavior*, vol. 32, pp. 233–242, 1984.

[7] W. J. Sutherland, "Aggregation and the ideal free distribution," *Journal of Animal Ecology*, vol. 52, pp. 821–828, 1983.

[8] J. Finke, K. Passino, and A. Sparks, "Cooperative control via task load balancing for networked uninhabited autonomous vehicles," in *Proceedings of the IEEE Conference on Decision and Control*, (Maui, Hawaii), pp. 31–36, December 2003.

[9] S. K. Subramanian and J. B. Cruz, "Adaptive models of pop-up threats for multi-agent persistent area denial," in *Proceedings of the IEEE Conference on Decision and Control*, (Maui, Hawaii), pp. 510–515, Dec. 2003.

[10] A. E. Gil, K. M. Passino, S. Ganapathy, and A. Sparks, "Cooperative scheduling of tasks for networked uninhabited autonomous vehicles," in *Proceedings of the IEEE Conference on Decision and Control*, (Maui, Hawaii), pp. 522–527, December 2003.

[11] E. Frazzoli and F. Bullo, "Decentralized algorithms for vehicle routing in a stochastic time-varying environment," in *43rd IEEE Conference on Decision and Control*, (Paradise Island, Bahamas), December 2004.

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

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

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