Proceedings of the 1st IFAC Workshop on Estimation and Control of Networked Systems, Venice, Italy, September 24-26, 2009

On distributed search in an uncertain environment

Matthias J. Feiler *

* *LGT Capital Management Ltd., Asset Allocation & Research, 8808 Pfäffikon, Switzerland (e-mail: mfeiler@ieee.org)*

Abstract: The paper investigates the case where N agents solve a complex search problem by communicating to each other their relative successes in solving the task. The problem consists in identifying a set of unknown points distributed in an n–dimensional space. The interaction rule causes the agents to organize themselves so that, asymptotically, each agent converges to a different point. The emphasis of this paper is on analyzing the collective dynamics resulting from *nonlinear* interactions and, in particular, to prove convergence of the search process.

1. INTRODUCTION

Research in the area of complex systems has increased substantially over the past years. The term is used in many different disciplines to denote the collective dynamic behavior of a large group of individuals (agents) each following his or her own (local) rules. A recurrent theme is the *emergence* of higher–order features that result from the interaction of such individuals. The interaction leads to a dynamics that is substantially different and cannot be guessed from the behavior of a single agent. Phase transitions and their determinants are studied and used to explain sudden disruptions in complex interconnected systems such as stock markets, traffic, social and political systems, see [1] for an overview and e.g. [2].

An important strand of literature has been created around the notion of flocking where agents interact in order to achieve a common group objective. Examples of such a behavior can be found in various forms of animal grouping (bird formation flying, fish schools, swarms of bees), see [3], [4]. Flocking dynamics has essentially three characteristics [5]: *cohesion*, i.e. stay nearby flockmates, *collision avoidance* and *alignment*, i.e. attempt to match the velocity vector of flockmates. Different aspects of these elementary rules have been studied theoretically, in particular the emergence of a common alignment [6]. A simple interaction rule states that if every agent adjusts its orientation according to the average of its nearest neighbors then all agents converge to a common orientation. Conditions on the emergence of such a *consensus* even in the presence of a switching network topology have been derived in [7].

In statistical mechanics, the Ising model provides a framework for the study of collective phenomena that can be attributed to local imitation of neighboring agents. A large variety of socio-economic processes that involve collective aggregate decision making has been analyzed using this model, for example the price dynamics of assets traded in the financial market [8] or the dynamics of opinions [9]. In the *minority game*, agents compete for minority membership, a somehow paradoxical situation in which a global equilibrium is impossible. Every agent is given a binary choice of actions (say go or stay, see the original problem proposed in [10]) where the payoff is positive only if the action corresponds to the choice of the minority. If all agents believe the others will go, nobody will go, thus invalidating the belief. The game has been used as a prototypical model of financial markets in which profit depends on the fact that an agents buys when the majority sells and vice versa [11].

The problem considered in this paper is also of a competitive nature. In our case, agents compete for N pieces of information represented as points in n– dimensional space. The information is revealed one piece at a time where the sequence is random and infinite (i.e. all N points appear infinitely often). The problem arises naturally in the identification of time–varying systems and was originally posed by Narendra, Feiler in 2003 [13], [14]: The elements of a vector $\theta \in \mathbb{R}^n$ represent the unknown parameters of a dynamical system. θ can assume one of a finite number N of constant values $\theta_i,\ i \in \Omega = \{1, 2, \dots, N\}$ at random instants of time t. For example, the time-variation may be governed by an ergodic Markov–chain and there exists an interval of finite length within which every $i \in \Omega$ is assumed at least once. N estimation models are set up with the objective of identifying the N values. The identification procedure consists in approaching the models $\hat{\theta}$ towards the unknown parameters θ such that, ultimately, the two sets coincide.

The work following the above problem statement was primarily focused on deriving conditions for the stability of an adaptive controller based on multiple identification models [15], [16], [17]. In this paper, we study the convergence properties of the identification procedure itself. To this end, we replace the identification models by simple search agents $\hat{\theta}_i$ that freely move in parameter space with the objective of finding the unknown positions θ_i representing some useful information. Once a piece of information is found it cannot be exploited by any other agent. Also, the agents have bounded capacity and can absorb only one piece of information. Thus an agent has not necessarily an interest of "being there first" but of finding a *different* piece of information than all the other agents. Here we allow for a global equilibrium by assuming that there are N agents and N pieces of information to be found. We are interesting in the collective dynamics leading towards that equilibrium.

In section II we present the search algorithm and state the convergence problem in mathematical terms. Section III contains a detailed analysis of the algorithm, in particular its properties as a dynamical system as well as the proof of convergence. In Section IV we conclude with simulation results.

2. SEARCH USING MULTIPLE AGENTS

Let us refer to the set $\mathcal{S} = \{\theta_1, \theta_2, \ldots, \theta_N\} \subset \mathbb{R}^n$ as *information set* and $\mathcal{A} = \{\hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_N\} \subset \mathbb{R}^n$ as the *agent set*. For an agent $\hat{\theta}$ to find a piece of information θ means that $\hat{\theta}(t) \to \theta$ as $t \to \infty$.

2.1 Problem Statement

Using the above variables, the problem can be stated succinctly as follows: Determine an algorithm such that for every $\theta \in \mathcal{S}$ there is a $\hat{\theta} \in \mathcal{A}$ such that $\|\hat{\theta}(t) - \theta\| \to 0$ as $t \to \infty$. In other words, for every piece of information there exists one agent that asymptotically finds it.

Remark: The implicit assumption made here is that N is known and, hence, the appropriate number of agents can be determined. The assumption can be relaxed to read $M \geq N$ where M is the number of models.

There are two situations where the problem is (almost) trivial:

(i) The index $j \in \Omega$ of the information revealed at instant of time t is known for every $t > 0$. In such a case a simple rule may be used to determine which agent looks for which information (e.g. the one that carries the same index). Only a single agent would be active at every instant (knowing that the information corresponding to its index is being revealed).
(ii) Suppose the agent and information sets are known to be *aligned*, i.e. for each $\theta \in S$ there exists a unique closest (when compared to the other agents) $\hat{\theta} \in \mathcal{A}$. Again, only a single agent moves at every instant of time (the one that is closest).

The difficulty of the search problem considered in this paper is that both the revelation of information over time as well as the location of information in space is random. We refer to such a situation as an *uncertain environment*.

2.2 The algorithm

The algorithm was first proposed in [13], motivated by the work of Haruno et al. [12] and based on the idea that all agent search simultaneously while communicating to each other their successes in finding a given piece of information $\theta \in S$. We define the error vector

$$\tilde{\theta}_i(t,t) = \hat{\theta}_i(t) - \theta(t) \tag{1}$$

Notice that the first time argument $\tilde{\theta}_i(t,\cdot)$ refers to the evolution of the search process while the second argument $\tilde{\theta}_i(\cdot,t)$ refers to the revelation of information $\theta(t)$ over time. The two processes may evolve on different time-scales. To make the search problem meaningful we assume that new information appears at a slower rate than the time needed to update an agent's "knowledge". Also, the fact that the total information is finite and that the same information appears over and over again makes the search problem meaningful. The error dynamics for agent i's search is given by

$$\tilde{\theta}_i(t+1,t) = \Big(1 - \eta_i[\tilde{\theta}(t,t)]\Big)\, \tilde{\theta}_i(t,t) \tag{2}$$

The central idea of the algorithm is contained in the term η_i which determines the step–size by which agent i is being updated.

$$\eta_i[\tilde{\theta}(t,t)] = \frac{\|\tilde{\theta}_i(t,t)\|^{-2}}{\sum_{k=1}^{N} \|\tilde{\theta}_k(t,t)\|^{-2}} \tag{3}$$

Equation (3) defines a communication protocol among the agents as follows. The step by which an agent "'moves" towards a given point, say θ_j depends both on its distance $\|\tilde{\theta}_i(t,t)\|$ and the distances of all the other agents. If an agent is already close to θ_j relative to the other agents, it is rewarded by a large step–size. The same agent obtains a smaller step–size if also the other agents are close, i.e. when the denominator in (3) is large. Notice also, that at every instant t all agents are active regardless of their relative positions. In other words, the protocol does not follow a "winner takes it all"– policy but assigns a non–zero step–size to every agent as long as the information has not been absorbed by any agent. This rule prevents the agents from getting locked, in particular if there is one agent closer to more than one $\theta_j \in S$ than all the others. The winner–rule would make it impossible for that agent to decide which piece of information to take. We observe that

$$\begin{aligned} \eta_i[\tilde{\theta}] &\to 1 \quad \text{if } \tilde{\theta}_i \to 0 \\ \eta_i[\tilde{\theta}] &\to 0 \quad \text{if } \tilde{\theta}_{j\neq i} \to 0 \end{aligned} \tag{4}$$

The protocol (3) ensures that no two models converge to the same element.

A final remark regards the nature of the dynamical system associated with the algorithm defined by (2), (3). The presence of the (stochastic) process governing the revelation of information makes the system non–autonomous. This means that at an equilibrium point we have

$$\tilde{\theta}_i^*(t+1,t) = \tilde{\theta}_i^*(t,t) \quad \text{for all } t > 0 \tag{5}$$

2.3 The equilibrium set

From equations (2) and (5) it is clear that for $\tilde{\theta}_i^*(t,t)$ to be an equilibrium point the product $\eta_i[\tilde{\theta}_i^*(t,t)]\,\tilde{\theta}_i^*(t,t) \equiv 0$ for all $t > 0$. It is also clear, that neither one of the factors can be zero for all $t > 0$ since convergence to $\theta_j \in S$ implies not to converge to another element $\theta_{k\neq j} \in S$ that is distinct from θ_j. At the equilibrium, there will be another agent at θ_k. But this implies that if indeed convergence takes place,

Lemma 1. Convergence takes place simultaneously over all elements in S.

Proof: Suppose all but one piece of information in S, say θ_k, has been absorbed by (at least) one agent. Then $\eta_i[\theta(t_k,t_k)] > 0$ for all $i \in \Omega$ where t_k are the instants of time for which $\theta(t_k) \equiv \theta_k \in S$. This means that equation 5 can only be satisfied for all $t > 0$ if there is an agent for *every* element in S. ■

The equilibrium set is defined as

$$I = \{\tilde{\theta}_i(t,t)\,|\,\tilde{\theta}_{f(i)}(t,t) = 0 \text{ if } \theta(t) \equiv \theta_{f(i)},\ i \in \Omega\} \tag{6}$$

where $f(\cdot) : \Omega \to \Omega$ is a rearrangement of the agent indices such that every agent has the same index as the point to which it converges. Note that there are $N!$ possible rearrangements of the agent indices. It is not clear which agent converges to which element of the set S.

Remark: If the number of agents M exceeds that of the pieces of information N there will be $M - N$ agents converging at positions not contained in the set S.

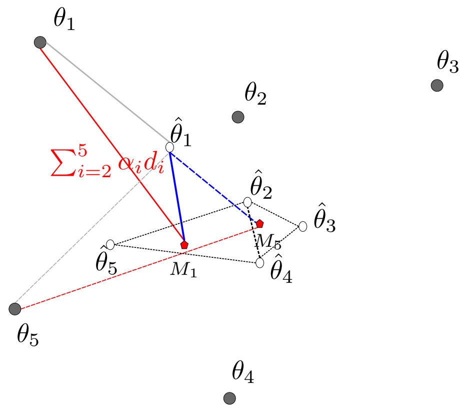

Fig. 1. Proof of theorem 1

3. ANALYSIS

The following notation is used:

$$d_i(\theta) = \|\hat{\theta}_i - \theta\| = \|\tilde{\theta}_i\| \tag{7}$$

denotes the distance of an agent to the point θ. The explicit dependence on time t is omitted wherever appropriate to simplify notation. We use the superscript $^+$ to denote the value of a variable after the algorithm has been applied once. The length of an update step is denoted as

$$s_i(\theta) = d_i - d_i^+ = \eta_i \|\tilde{\theta}_i\| \tag{8}$$

Without loss of generality, the index $i = 1$ is used to denote a particular element of the sets $\mathcal{S}$ or $\mathcal{A}$ as opposed to the other elements of the set $i = 2, \ldots N$. As stated above, the index set is denoted as $\Omega = \{1, 2, \ldots, N\}$.

3.1 Convergence to a Convex Hull

In order to motivate the problem, we assumed that the agents $\mathcal{A}$ are initialized far from the points $\mathcal{S}$. Hence, none of the agents is distinguished with respect to any point and we expect all agents to behave similarly. Let $H(\mathcal{S})$ be the convex hull of the *information set* $\mathcal{S}$.

Lemma 2. Given $i \in \Omega$ with $\hat{\theta}_i \notin H(\mathcal{S})$, then $\hat{\theta}_i$ converges monotonically to $H(\mathcal{S})$, i.e. $\text{dist}(\hat{\theta}_i(t), H(\mathcal{S})) := d_{iH(\mathcal{S})}(t) \to 0$ monotonically as $t \to \infty$.

Proof: From the algorithm we see that every agent reduces its distance to an element contained in $H(\mathcal{S})$ at every instant $t > 0$. The lemma follows from the definition of the distance of an agent $\hat{\theta}_i$ to the set $H(\mathcal{S})$

$$d_{iH(\mathcal{S})}(k) := \min_{\xi \in H(\mathcal{S})} \|\hat{\theta}_i(k) - \xi\|$$

■

Hence, without loss of generality, all agents can be initialized inside the convex hull.

3.2 Ordering of the steps

The amount by which the agents reduce their distances to the prevailing point $\theta \in S$ is ordered according to their relative distances from that parameter. Using the definition of η_i from equation (3) we obtain

$$s_i(\theta) = \eta_i d_i = \frac{d_1^2 \ldots d_{[i-1]}^2\, d_i\, d_{[i+1]}^2 \ldots d_N^2}{\sum_{k=1}^{N} \prod_{l \neq k} d_l^2} \tag{9}$$

Let the indices $i \in \Omega$ be assigned such that

$$d_1 < d_2 < \cdots < d_N \tag{10}$$

The following lemma regarding the ordering of the steps holds:

Lemma 3.

$$d_i < d_{i+1} \quad \Rightarrow \quad \eta_i d_i > \eta_{i+1} d_{i+1} \qquad \forall\, i \in \Omega \tag{11}$$

Proof: From equation (9) we have:

$$\begin{aligned} s_i &= \frac{1}{d_S}\, d_1^2 \ldots d_{[i-1]}^2\, d_i\, d_{[i+1]}^2 \ldots d_N^2 \\ s_{i+1} &= \frac{1}{d_S}\, d_1^2 \ldots d_i^2\, d_{[i+1]}\, d_{[i+2]}^2 \ldots d_N^2 \end{aligned} \tag{12}$$

where $d_S = \sum_{k=1}^{N} \prod_{l \neq k} d_l^2$. We obtain

$$\frac{s_i}{s_{i+1}} = \frac{d_{i+1}}{d_i} > 1 \quad \text{since } d_i < d_{i+1}$$

3.3 Agent dispersion

The dispersion between two agents is simply the Euclidean distance $\|\hat{\theta}_1 - \hat{\theta}\|$ in $\mathbb{R}^n$. Here $\hat{\theta}_1$ refers to agent 1 and $\hat{\theta}$ to any other agent. We now relate this distance to a point $\theta \in \mathcal{S}$ and define

$$\rho(\theta) = \big|\|\hat{\theta}_1 - \theta\| - \|\hat{\theta} - \theta\|\big| \tag{13}$$

Using (7) we write

$$\rho(\theta) = |d_1(\theta) - d(\theta)| \tag{14}$$

$\rho(\theta)$ is a *relative dispersion* as seen from the perspective of θ. From the triangle inequality we obtain that

$$\rho(\theta) \leq \|\hat{\theta}_1 - M\| \tag{15}$$

We separate $\hat{\theta}_1$ from the rest of the group and let $\mathcal{A}'$ denote that rest $\mathcal{A}' = \{\hat{\theta}_2, \ldots, \hat{\theta}_N\}$. The distance of the agents from θ is given by $d_2, \ldots, d_N$ respectively. A convex combination reads

$$d_{\mathcal{A}'} = \sum_{i=2}^{N} \alpha_i d_i \quad \text{where} \ \sum_{i=2}^{N} \alpha_i = 1,\ \alpha_i > 0\ \forall\, i. \tag{16}$$

From the geometry it is clear that any ball of radius $d_{\mathcal{A}'}$ has a non–zero intersection with the convex hull of $\mathcal{A}'$. In other words, there exists a point $M \in H(\mathcal{A}')$ which is at a distance $d_{\mathcal{A}'}$ from θ. M may be interpreted as an *average* position of the agents in $\mathcal{A}'$. If $\alpha_i = 0$ for some $i \in \{2, \ldots, N\}$ then M may be a boundary point of $H(\mathcal{A}')$. The generic situation, however, is that M is inside $H(\mathcal{A}')$. This is displayed in figure 1.

Lemma 4. For any θ there exists a point $M \in H(\mathcal{A}')$ such that

$$\rho(\theta)^+ > \rho(\theta) \tag{17}$$

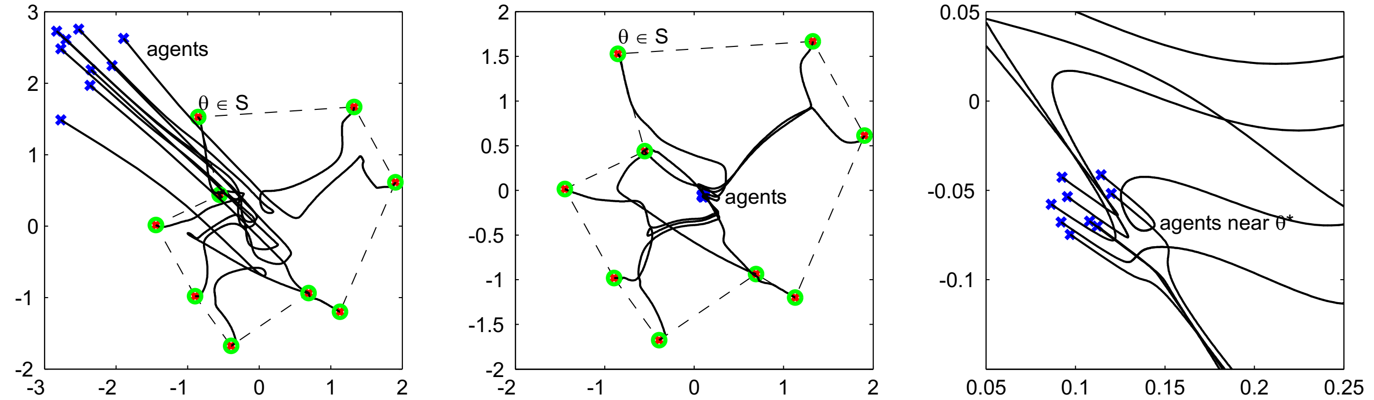

Fig. 2. Convergence of the search process. *Left:* all agents are initialized far from the points. *Middle:* Agents are initialized in a neighborhood of the saddle point θ^*. *Right:* Detailed view of the first iterations of the algorihtm started at θ^*

Proof: For equation (17) to hold, the following must be satisfied

$$\begin{array}{ll} s_M - s_1 > 0 & \text{if} \quad d_1 > d_M \\ s_1 - s_M > 0 & \text{if} \quad d_M > d_1 \end{array} \tag{18}$$

The conditions are formulated in terms of the first agent $i = 1$ but are the same for any $i \in \Omega$. Notice that $s_M = \sum_{i=2}^{N} s_i$. Obviously $s_M \neq \eta_M d_M$ since M is not a *real* agent to which the algorithm can be applied. Let us first consider $d_1 > d_M$. We define

$$\begin{array}{ll} \alpha_i^* > 0 & \text{if } d_1 > d_i \\ \alpha_i^* = 0 & \text{if } d_1 \leq d_i \end{array} \tag{19}$$

and such that $\sum_{i=2}^{N} \alpha_i^* = 1$. From the above arguments we know that there exists an M in $H(\mathcal{A}')$ such that $d_M = \sum_{i=2}^{N} \alpha_i^* d_i$.

$$s_M = \sum_{i=2}^{N} \alpha_i^* \eta_i d_i > \sum_{i=2}^{N} \alpha_i^* \eta_1 d_1 = s_1 \tag{20}$$

since by lemma 3

$$\eta_i d_i > \eta_1 d_1 \text{ whenever } d_1 > d_i \tag{21}$$

The same idea can be applied to the second inequality in (18). ■

We are now ready to state our main result.

Theorem 1. For every $\theta \in \mathcal{S}$ there exists a $\hat{\theta} \in \mathcal{A}$ such that

$$\lim_{t \to \infty} \hat{\theta}(t) = \theta \quad a.s. \tag{22}$$

under the dynamics defined by the algorithm in equations (2) and (3).

Remark: $\hat{\theta}(t)$ is a stochastic variable whose evolution is defined by the algorithm which, in turn, is driven by the stochastic force $\theta(t)$[1]. Hence, the best we can obtain is convergence with probability 1.

Lemma 4 states that for any $\theta \in \mathcal{S}$, say θ_1, there exists an $M \in H(\mathcal{A}')$ such that $\rho_{t+1}(\theta_1) > \rho_t(\theta_1)$ where $t > 0$ are instants of time. The central question is whether this inequality can be made invalid in any of the subsequent iterations of the algorithm[2]. When considering the dynamics of the agents it is useful to introduce the notion of a cycle.

Definition 1. Let $T > 0$ be such that $\theta(t + T) = \theta(t)$ and $\theta(t + t_i) = \theta_i$ for some $t_i \in [0 \;\; T]$ and every $i \in \Omega$. T is called a *cycle*.

While it cannot be excluded that an agent returns to its original position after a cycle $\hat{\theta}_1(t + T) = \hat{\theta}_1(t)$ the same is not true for the distance between $\hat{\theta}_1$ and a point M in the convex set $H(\mathcal{A}')$.

Lemma 5. For every $\theta \in S$, say θ_1 there exists a subsequence $\{t_s\}_{s \geq 0}$ of $\{t\}_{t \geq 0}$ such that

- $t_s \to \infty$ as $s \to \infty$.
- $\theta(t_s) = \theta_1$ for all $s \geq 0$.
- and the following holds:

$$\rho_{t_{s+1}}(\theta_1) > \rho_{t_s}(\theta_1). \tag{23}$$

(Qualitative) Proof. If we track the motion of $\hat{\theta}_1$, we know from lemma 4 that $d_1(\theta) - d_M(\theta)$ increases whenever $d_1(\theta) > d_M(\theta)$ and vice versa. This holds at every step t, where $t_s \leq t \leq t_{s+1}$ and independent of θ. It follows that the diameter (defined as the smallest Euclidean distance among two elements) of the set $\hat{\theta}_1 \cup \mathcal{A}'$ increases. For if we suppose in contradiction that it decreases, any increase must be "undone" by a (net) step towards an opposite point. But during such a motion, the reverse relation $d_1(\theta) < d_M(\theta)$ holds and $|d_1(\theta) - d_M(\theta)|$ increases further. But this means in particular that at the end of a cycle M can be chosen in such a way that inequality (23) holds. □

Proof of the Theorem 1. We first investigate the convergence from the perspective of θ_1. To this end, define a sequence $\{M_{t_s}\}_{s \geq 0}$ according to lemma 5. The corresponding sequence $\{\rho_{t_s}\}_{t_s \geq 0}$ is monotonically increasing. Since both M_{t_s} and $\hat{\theta}_1$ are contained in the bounded region $H(\mathcal{S})$, ρ_{t_s} is bounded above by the largest diagonal $D_1 = \max_{i \in \Omega / \{1\}} \|\theta_i - \theta_1\|$ in the set $\mathcal{S}$. This means that $\rho_{t_s} \to D_1$ as $s \to \infty$ which is equivalent to

$$\lim_{s \to \infty} \hat{\theta}_i(t_s) = \theta_1 \quad \text{for some } i \in \Omega \tag{24}$$

The same argument holds with respect to every other point $\theta_i \in S$. ■

Remark: We do not know which agent converges to θ_1 even though ρ_{t_s} is defined in terms of $\hat{\theta}_1$. What we rely on in the

[1] If we solve equation (2) for $\hat{\theta}_i(t+1)$ using equation (1) the driving force appears in the term $\eta_i[\tilde{\theta}(t,t)]$.

[2] It is clear that this is impossible if $\theta(t) \equiv \theta_1$ for all $t > 0$. If, on the other hand, $\theta(t)$ switches to a new value a new "geometry" is in place and lemma 4 holds with respect to a new point $\theta_{i \neq 1}$.

proof is the monotonic increase (over a subsequence) of the dispersion among the agents, i.e. the distance between $\hat{\theta}_1$ and a point M in $H(\mathcal{A}')$. M may correspond to an element of $\mathcal{A}'$ in which case it is a candidate for the convergence to θ_1.

4. DISCUSSION

Figure 2 displays the results of a simulation involving $N = 9$ agents. In the first experiment, the agents are initialized close to each other but far from the points. The monotonic convergence stated in lemma 2 appears very distinctly. Once the agents have entered a convex hull of $\mathcal{S}$ the dynamics of the agents is much more convoluted. This is also observed in the second experiment, where the agents are initialized inside the convex hull. A qualitative feature observed in many simulations is that the agents seem to gather at a point inside $H(\mathcal{S})$ before fanning out and converging to the different points contained in S. At this point it is not evident which agent converges to which point. Also, the way the convergence takes place is virtually invisible to the naked eye. In order to structure the observations that can be made in simulations we introduce a simplification which helps us locate the "gathering point" of all the agents.

First, we assume that the time–variation is periodic. Second, instead of updating the agents at every instant of time, their location is fixed. The adjustment vectors which result at every instant are summed up and used to update the location of the model only after the cycle is complete. Last, all agents are initialized at the same point. In this case no agent has a relative advantage, and the communication (of relative successes) becomes obsolete. The step-size equation (3) reads:

$$\eta_i(t) = \frac{1}{N} \tag{25}$$

In this special set–up, the only equilibrium point is at

$$\theta^* = \frac{1}{N}\sum_{i=1}^{N}\theta_i \tag{26}$$

and it is easily seen that θ^* is stable and globally attractive. The point helps us distinguish to phases in the convergence process: In the initial stage, the models merely "position" themselves (and approach a neighborhood of θ^*) without noticeably increasing their relative positions (agent dispersion). In the second stage, the agents diverge from θ^* while approaching the points of the set S. θ^* has a remarkable influence on the evolution of agents even though it is not an equilibrium point for the original dynamics.

This suggests the existence of a hyperbolic invariant set, e.g. a limit cycle that is reached when all agents start from identical positions. The stable manifold of this set is followed in the first stage of the convergence process while the unstable manifold dominates the dynamics in the second stage. The interesting point to note here is that the hyperbolic structure is the result of the competition introduced by the communication protocol defined in equation (3).

REFERENCES

[1] Tesfatsion, L.: *Agent-based computational economics: modeling economies as complex adaptive systems.* Information Sciences, Volume 149, Issue 4, February 2003.

[2] Mauboussin, M.J.: *Revisiting market efficiency: the stock market as a complex adaptive system.* Journal of Applied Corp. Finance, Volume 4, April 2005.

[3] Shaw, E.: *Fish in schools.* Natural History, vol. 84, no.8, 1975.

[4] Toner, J.; Tu, Y.: *Flocks, herds, and schools: A quantitative theory of flocking.* Physical Review E, vol. 58, no. 4, Oct. 1998.

[5] Reynolds, C.W.: *Flocks, herds, and schools: a distributed behavioral model.* Computer Graphics, vol. 21, no. 4, July 1987.

[6] Olfati–Saber, R.: *Flocking for Multi-Agent Dynamic Systems: Algorithms and Theory.* IEEE Transactions on Automatic Control, vol. 51, no. 3, March 2006.

[7] Jadbabaie, A.; Lin, J.; Morse, A.S.: *Coordination of groups of mobile autonomous agents using nearest neighbor rules*, IEEE Transactions on Automatic Control, vol. 48, no. 6, 2003.

[8] Sornette, D.; Zhou, W.X.: *Importance of Positive Feedbacks and Over–Confidence in a Self–Fulfilling Ising Model of Financial Markets*, Journal of Economic Behavior and Organization,

[9] Sznajd–Weron, K.; Sznajd, J.: *Opinion Evolution in Closed Community*, International Journal of Modern Physics C, vol. 11, no. 6, 2000.

[10] Arthur, W.B.: *Inductive Reasoning and Bounded Rationality*, American Economic Assoc. Papers and Proc. 84, 1994.

[11] Challet, D.; Marsili, M.; Zhang, Y.C.: *Minority Games: Interacting Agents in Financial Markets*, Oxford University Press, Inc. New York, USA, 2005.

[12] Haruno, M.; Wolpert, D.M. and Kawato M.: *MOSAIC Model for Sensorimotor Learning and Control.* Neural Computation, 13: 2201–2220, 2001.

[13] Narendra, K.S.; Driollet, O.A.; Feiler, M.J.; Koshy, G.:*Adaptive Control Using Multiple Models, Switching and Tuning.* International Journal of Adaptive Control and Signal Processing, vol. 17, pp.87-102, 2003.

[14] Feiler, M.J.; Narendra, K.S.: *Adaptive Control of Rapidly Time-Varying Systems.* In: Proceedings of the 12th Yale Workshop on Adaptive and Learning Systems, p. 33-42, Yale University, New Haven, CT, USA, 2003.

[15] Feiler, M.J.; Narendra, K.S.: *Simultaneous Identification and Control of Linear Time–Varying Plants using Multiple Models.* Technical Report No. 0401, Center for Systems Science, Yale University, New Haven, CT, USA, March 2004.

[16] Feiler, M.J; Narendra, K.S.:*Simultaneous Identification and Control of Time–Varying Systems* In: Proceedings of the 45th IEEE Conference on Decision and Control, 2006.

[17] Feiler, M.J; Narendra, K.S.:*Identification and Control Using Multiple Models* In: Proceedings of the 14th Yale Workshop on Adaptive and Learning Systems, p. 33-42, Yale University, USA, 2008.