Density classification with a continuous-time binary consensus protocol with hysteretic units

Marco A. Montes de Oca and Louis F. Rossi

Dept. of Mathematical Sciences, University of Delaware, Newark, DE 19716
{mmontes, rossi}@math.udel.edu

Abstract

In this paper, we present a continuous-time binary consensus protocol whereby entities connected via a directed ring topology solve the one-dimensional density classification problem. In our model, entities behave as non-ideal relays, that is, they have memory of the trajectory of an internal state variable, which gives them hysteretic properties. We show that this feature is necessary for the system to reach consensus on the state shared by the initial majority.

Introduction

The density classification problem (also known as the majority problem), consists in classifying finite linear binary strings according to whether they have a majority of 0's or 1's. This problem has been the subject of numerous studies in the cellular automata literature (see, e.g., Mitchell et al. (1994); Land and Belew (1995); Fukś (1997, 2002); Alonso-Sanz and Bull (2009); Fatès (2013)) because it has a simple formulation and illustrates very well the idea of "emergent computation" since the cells interact only locally and do not have access to the global structure they are trying to classify. In a cellular automata setting, the density classification problem translates into finding evolution rules for cells that make the automata be all equal to 0 if initially there were more than half of the cells in a 0 state, or 1 if initially there were more than half of the cells in a 1 state.

The idea of emergent computation is also present in another computational paradigm called swarm intelligence (Bonabeau et al., 1999; Dorigo and Birattari, 2007). In this paradigm, relatively simple agents locally interact with one another and with their environment to produce self-organized spatio-temporal patterns that represent solutions to problems that no individual agent could solve on its own (e.g., finding shortest paths (Goss et al., 1989), sorting (Deneubourg et al., 1990), or constructing nests (Grassé, 1959)).

As we explain in the next section, the density classification problem is relevant in swarm intelligence because some of the methods that solve it may be used as collective decision-making mechanisms for swarms. In this paper, we further explore the connection between cellular automata and swarm intelligence by presenting and analyzing a consensus protocol[1] on networks of agents derived from previous work on collective decision-making in swarms (Montes de Oca et al., 2012).

Our consensus protocol may be seen as encoding evolution rules for continuous-time cellular automata with memory. It may also be thought of as model of social influence in a group whose members observe the actions performed by other individuals, increasing, as a result, their tendency to perform the observed actions. We are interested in this type of mechanisms because we want to eventually endow swarms of robots or agents with collective-decision mechanisms that are robust, flexible and effective in real environments. Our approach is backed by recent studies on social learning (Rendell et al., 2010), that support the idea that learning from the observation of others' actions is a mechanism whereby individuals indirectly probe the environment. In a swarm, which is typically composed of many individuals, using the behavior of others as a guide provides each individual with potentially many indirect channels for obtaining information about their environment, and thereby increasing the amount of information they use to make decisions.

The key finding presented in this paper is that if agents influence each other as if arranged in a unidirectional ring, and each agent integrates over time information coming to it from its neighbor using a mechanism akin to exponential smoothing (Gardner Jr., 2006), then symmetric blocks of 0's or 1's propagate through the network indefinitely. Moreover, we provide evidence that when the symmetry of these blocks is broken (that is, there is a majority of 0's or 1's), then the information wave propagates for a finite amount of time and eventually dies out, which translates into the population of agents reaching a consensus. Finally, the state on which the population reaches a consensus corresponds to that of the initial majority. In other words, we provide evidence that

[1] We use the term *protocol* to comply with literature tradition in communication networks and where interaction rules are called protocols. See, for example, (Mesbahi and Egerstedt, 2010).

the proposed continuous-time consensus protocol with hysteretic units solves the density classification problem. Ongoing work is aimed at analytically determining whether our observations hold for strings of any length, and for any distribution of 0's and 1's within these strings. And if so, how much time is needed for the system to converge.

Colective Decision-Making in Swarms

In (Montes de Oca et al., 2012), we proposed a social influence model whose dynamics can be used as a collective decision-making mechanism for swarms of robots that need to collectively choose the most efficient of two alternative actions (henceforth referred to as *Decision-Making Model or DM model*). In the DM model, each of a set of n agents can be in one of two states (represented with a binary variable $X_i \in \{0, 1\}$, with $i = 1, 2, \ldots, n$). In applications of the DM model, an agent's state can represent, for example, a robot's preferred action or current belief of the state of an environmental variable.

The DM model is a discrete-time model where at each time step t of the system's evolution, an agent i might be able to observe the state of another random agent $j \neq i$. When agent i observes the state of another agent j, the observing agent i updates an internal real-valued variable S_i, which we call tendency, as follows:

$$S_i^{t+1} = (1 - \alpha)S_i^t + \alpha X_j^t \,, \tag{1}$$

where $0 \leq \alpha \leq 1$ determines the relative weight given to the agent's latest observation (X_j^t) and the agent's accumulated experience (S_i^t).

After updating its tendency, an agent updates its state as follows:

$$X_i^{t+1} = \begin{cases} 1, & \text{if } S_i^{t+1} \geq \lambda \\ 0, & \text{if } S_i^{t+1} \leq \mu \\ X_i^t, & \text{if } \mu < S_i^{t+1} < \lambda \,, \end{cases} \tag{2}$$

where $\mu + \lambda = 1$ (the reason for this constraint will become apparent later). Eq. 2 implements a sort of dynamic memory that allows the agent to integrate its observations over time. By properly choosing values for the parameters α, μ, λ, and the initial conditions X_i^0 and S_i^0, one can control the imitation behavior of agent i. While in principle, each agent may have different values for its parameters, in the DM model, α, λ and μ, are constant constant and common to all agents. An example of the behavior of an individual agent in the DM model is shown in Figure 1.

In the DM model the population is reshuffled randomly so that each individual observes a different agent at each time step. Additionally, one single agent may potentially influence more than one other agent in the group. The DM model's collective dynamics make the population reach a consensus on the state that at time step 0 is shared by most (that is, the majority) of the population. In (Montes de Oca

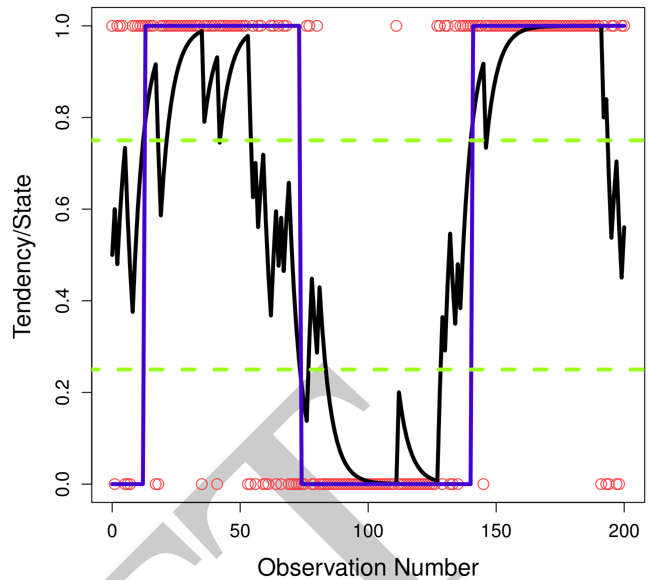

Figure 1: Single agent behavior in the DM model. Starting with an initialization of $S^0 = 0.5$ and $X^0 = 0$, an agent observes a stream of state values plotted as dots with values 0 or 1. The black line shows the evolution of the agent's tendency and the blue line shows the evolution of the agent's state. In this simulation, $\alpha = 0.2$, $\lambda = 0.75$ and $\mu = 0.25$ (shown as dotted lines).

et al., 2012), and (Montes de Oca et al., 2011), we show how this behavior can be used for optimal collective decision-making in robot swarms.

The DM model may be seen as a method to solve the density classification problem if its definition is relaxed. In particular, if the agents (cells) are allowed to be reshuffled, then the DM may solve it. In this paper, we explore the question of whether it is possible to solve the original density classification problem with a variation of the DM model that does not require reshuffling. In the following sections, we present such a variation as well as theoretical and experimental results that make us believe that the question can be answered positively.

Continuous-Time Consensus Protocol with Hysteretic Units

The protocol that we propose in this paper, henceforth referred to as *Consensus with Hysteresis or CH*, is in its basic form the continuous-time equivalent of Eq. 1. However, in the CH protocol, the communication topology of the population does not change over time and each agent influences exactly one other agent. In the remainder of this paper, we assume that individuals are arranged in a directed ring topology with an agent i influenced by agent j, where $j = i + 1$ or $j = i - 1$ (agents "look" to their right or left, respectively).

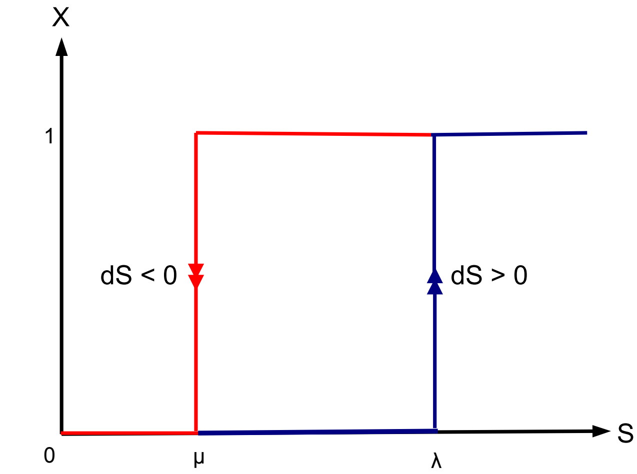

Figure 2: Hysteresis loop. Since hysteresis is nonreversible, we always assume that $dt > 0$. If $dS > 0$ (the input increases), the output evolves according to $\frac{dX}{dt} = g_1(S, X)$ (blue line). If $dS < 0$ (the input decreases), the output evolves according to $\frac{dX}{dt} = g_2(S, X)$ (red line).

In the CH protocol, the tendency-update equation for an agent i influenced by agent j ($i \neq j$) is:

$$\frac{dS_i}{dt} = \alpha(X_j - S_i)\,, \tag{3}$$

where S_i is now a continuous-time function and X_j is *not* a function but a nonlinear operator that defines a nonideal relay on S_j, that is, $X_j = R_{\mu,\lambda}(S_j)$, where $R_{\mu,\lambda}$ is a hysteresis operator with thresholds μ and λ (see Figures 2 and 3). In this work, we use the Duhem model of hysteresis (Visintin, 2006) to model this operator. The Duhem model defines X_j as the solution of the initial value problem:

$$\begin{cases} \dfrac{dX_j}{dt} = g_1(S_j, X_j)\left(\dfrac{dS_j}{dt}\right)^+ - g_2(S_j, X_j)\left(\dfrac{dS_j}{dt}\right)^-, \\ X_j(0) = X_j^0, \end{cases} \tag{4}$$

where $x^+ = (|x|+x)/2$ and $x^- = (|x|-x)/2$ for any $x \in \mathbb{R}$. Additionally, g_1 and g_2 are given nonnegative functions that represent the paths of evolution of the pair (S, X) for increasing and decreasing S, respectively (see Figure 2).

Rewriting Eq. 3 into standard form, we have:

$$\frac{dS_i}{dt} + \alpha S_i = \alpha X_j\,,$$

whose solution may be found by using the integrating factor $e^{\int \alpha dt} = e^{\alpha t}$. After multiplying and rearranging, we obtain

$$\frac{d}{dt}\left(e^{\alpha t} S_i\right) = e^{\alpha t} \alpha X_j$$

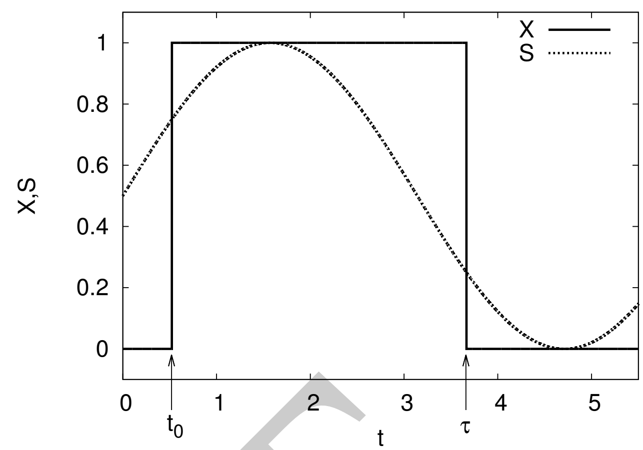

Figure 3: Hysteresis output of a nonideal relay operator $X = R_{\mu,\lambda}(S)$ with $\mu = 0.25$ and $\lambda = 0.75$.

Integrating both sides from t_0 to t:

$$\int_{t_0}^{t} \frac{d}{dw}\left(e^{\alpha w} S_i\right) dw = \int_{t_0}^{t} e^{\alpha w} \alpha X_j dw\,.$$

The right-hand side of this equation may be integrated by parts with $u' = \alpha e^{\alpha w}$, and $v = X_j$, which means $u = e^{\alpha w}$ and $v' = \frac{dX_j}{dw}$. After some substitutions and simplifications, we obtain:

$$\begin{aligned} S_i(t) = X_j(t) + e^{\alpha(t_0 - t)}(S_i(t_0) - X_j(t_0)) \\ - e^{-\alpha t} \int_{t_0}^{t} e^{\alpha w} \frac{dX_j}{dw} dw\,. \end{aligned} \tag{5}$$

To completely integrate Eq. 5, we must consider two cases: case 1 for $dS_j > 0$, and case 2 for $dS_j < 0$. For case 1, Eq. 5 becomes

$$\begin{aligned} S_i(t) = X_j(t) + e^{\alpha(t_0 - t)}(S_i(t_0) - X_j(t_0)) \\ - e^{-\alpha t} \int_{t_0}^{t} e^{\alpha w} g_1(S_j, X_j) dw\,. \end{aligned} \tag{6}$$

Similarly, for case 2, Eq. 5 becomes

$$\begin{aligned} S_i(t) = X_j(t) + e^{\alpha(t_0 - t)}(S_i(t_0) - X_j(t_0)) \\ - e^{-\alpha t} \int_{t_0}^{t} e^{\alpha w} g_2(S_j, X_j) dw\,. \end{aligned} \tag{7}$$

In our case, since the hysteresis loop may be seen as being composed of step functions, $\frac{dX_j}{dt} = g_1(S_j, X_j) = g_2(S_j, X_j) = 0$ in the interval $t \in (t_0, \tau)$, where t_0 is the time at which X last switched from 0 to 1 or vice versa, and τ is the time at which X next switches from 0 to 1 or vice versa (see Figure 3).

Table 1: Possible cases for the transition of states of agent i influenced by agent j.

Condition		Result	
X_i	X_j	dS_i	State transition
0	0	≤ 0	No
0	1	> 0	Yes, if $S_i < \lambda$
1	0	< 0	Yes, if $S_i > \mu$
1	1	≥ 0	No

Therefore, in the time interval (t_0, τ), the tendency of any agent i is given by

$$S_i(t) = X_j(t) + e^{\alpha(t_0 - t)}(S_i(t_0) - X_j(t_0)) \,. \quad (8)$$

Eq. 8 can be further simplified by recalling that in the interval (t_0, τ), X_j does not change, therefore

$$S_i(t) = (1 - e^{\alpha(t_0 - t)})X_j(t_0) + e^{\alpha(t_0 - t)}S_i(t_0) \,. \quad (9)$$

This last equation clearly shows that in intervals beyond (t_0, τ), $S_i(t)$ is actually a piecewise function that depends on how $X_j(t)$ evolves over time. Therefore, the CH protocol has a punctuated evolution, which allows us to analyze the dynamics of a system of agents governed by Eq. 9 by breaking time into intervals defined by a sequence of τ's. For example, we start with an interval $(t_0 = 0, \tau = \tau_1)$, then continue with $(t_0 = \tau_1, \tau = \tau_2)$, then with $(t_0 = \tau_2, \tau = \tau_3)$ and so forth, where $\tau_1, \tau_2, \tau_3, \ldots, \tau_k$ are the times at which some agent in the population switches from 0 to 1 or viceversa. Such an analysis is presented in the next section.

Consensus with Hysteretic Units: Analysis and Simulations

We focus first on the calculation of the sequence of τ's at which the system's states change as described above. In order to do this, we first notice that only a subset of all the possible initial conditions can lead to state transitions (see Table 1).

A state transition is possible only if the agents involved in an interaction have opposite states (this is an important observation because *it implies that consensus is an absorbing state*). In such cases, it is possible to calculate the time between state transitions τ.

In the second case of Table 1, $X_i(t_0) = 0$, $X_j(t_0) = 1$, and $S_i(t_0) < \lambda$, we have

$$S_i(\tau) = X_j(\tau) + e^{\alpha(t_0 - \tau)}\left(S_i(t_0) - X_j(t_0)\right) = \lambda \,.$$

Solving this equation for τ, we obtain

$$\tau = t_0 + \frac{1}{\alpha} \ln\left(\frac{X_j(t_0) - S_i(t_0)}{X_j(\tau) - \lambda}\right)$$

Since we are assuming that in the interval $(t_0, \tau]$, X_j does not change, then $X_j(\tau) = X_j(t_0) = 1$. Thus, we have

$$\tau = t_0 + \frac{1}{\alpha} \ln\left(\frac{1 - S_i(t_0)}{1 - \lambda}\right) \,. \quad (10)$$

Similarly, in the third case of Table 1, $X_i(t_0) = 1$, $X_j(t_0) = 0$, and $S_i(t_0) > \mu$, we have

$$S_i(\tau) = X_j(\tau) + e^{\alpha(t_0 - \tau)}\left(S_i(t_0) - X_j(t_0)\right) = \mu \,.$$

Solving this equation for τ, we obtain

$$\tau = t_0 + \frac{1}{\alpha} \ln\left(\frac{X_j(t_0) - S_i(t_0)}{X_j(\tau) - \mu}\right) \,.$$

In this case, if $X_j(\tau) = X_j(t_0) = 0$, then

$$\tau = t_0 + \frac{1}{\alpha} \ln\left(\frac{S_i(t_0)}{\mu}\right) \,. \quad (11)$$

From Eqs. 10 and 11, it is clear that the time to transition from 0 to 1 and from 1 to 0 is equal only when $S_i(t_0) = \frac{1}{2}$ and $\lambda + \mu = 1$ (this is the origin of the restriction in Eq. 2). This result is important in connection with the following observation.

That there is no state transition when $X_i = X_j$, does not mean that the tendency S_i does not change. In fact, if $X_i = X_j = 0$ and $S_i(t_0) > 0$, then $dS_i < 0$ and S_i decreases. Similarly, if $X_i = X_j = 1$ and $S_i(t_0) < 1$, then $dS_i > 0$ and S_i increases. This fact makes contiguous blocks of 0's or 1's particularly interesting because they make agents' tendencies go over or below the hysteresis thresholds and delay future state transitions (possibly assymetrically). Let us illustrate with two simple examples the situations that we can encounter as a consequence of having symmetric or asymmetric blocks of 0's or 1's.

Example 1: 0011 (symmetric blocks of 0's and 1's)

Imagine we initialize our units with the pattern `0011` (see Figure 4), that is, $X_1(0) = 0$, $X_2(0) = 0$, $X_3(0) = 1$, and $X_4(0) = 1$. Also, assume that $S_i(0) = 1/2$ for $i \in \{1, 2, 3, 4\}$. Table 2 shows the evolution of the system for the first 4 transitions.

This example shows that while the initial conditions for the agents' tendencies change over time, the symmetry of the system is kept and therefore, the state of the system is stable (more about this later). The duration of the inter-state transition intervals $\Delta\tau = \tau_{k+1} - \tau_k$ for some k is determined by the units that are at the interface between blocks of 0's and 1's. It is interesting to see the evolution of the duration of these intervals over time, especially for relatively large populations. In Figure 5, we show the duration of the inter-state transition intervals for a system with 10 units divided into 2 symmetric blocks of 0's and 1's.

Table 2: Punctuated evolution of a system of 4 units initialized with $S_i(0) = 1/2$ for $i \in \{1, 2, 3, 4\}$ and $X_1(0) = 0$, $X_2(0) = 0$, $X_3(0) = 1$, and $X_4(0) = 1$.

t	X_1	S_1	X_2	S_2	X_3	S_3	X_4	S_4
0	0	$\frac{1}{2}$	0	$\frac{1}{2}$	1	$\frac{1}{2}$	1	$\frac{1}{2}$
$\tau_1 = \frac{1}{\alpha}\ln(\frac{1}{2\mu})$	1	λ	0	μ	0	μ	1	λ
$\tau_2 = \tau_1 + \frac{1}{\alpha}\ln(\frac{\lambda}{\mu})$	1	$1 - \frac{\mu^2}{\lambda}$	1	λ	0	$\frac{\mu^2}{\lambda}$	0	μ
$\tau_3 = \tau_2 + \frac{1}{\alpha}\ln\left(\frac{1-\frac{\mu^2}{\lambda}}{\mu}\right)$	0	μ	1	$1 - \frac{\mu^2}{1-\frac{\mu^2}{\lambda}}$	1	λ	0	$\frac{\mu^2}{1-\frac{\mu^2}{\lambda}}$
$\tau_4 = \tau_3 + \frac{1}{\alpha}\ln\left(\frac{1-\frac{\mu^2}{1-\frac{\mu^2}{\lambda}}}{\mu}\right)$	0	$\frac{\mu^2}{1-\frac{\mu^2}{1-\frac{\mu^2}{\lambda}}}$	0	μ	1	$1 - \frac{\mu^2}{1-\frac{\mu^2}{1-\frac{\mu^2}{\lambda}}}$	1	λ

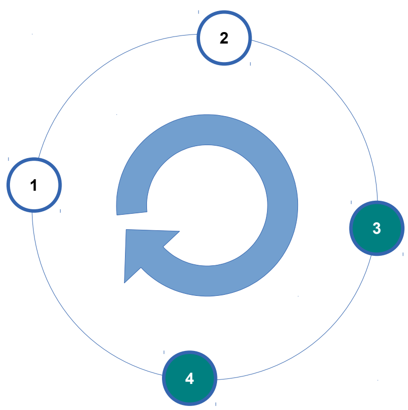

Figure 4: Example initialization of the CH protocol with four units. The direction of information flow is depicted with the big circular arrow.

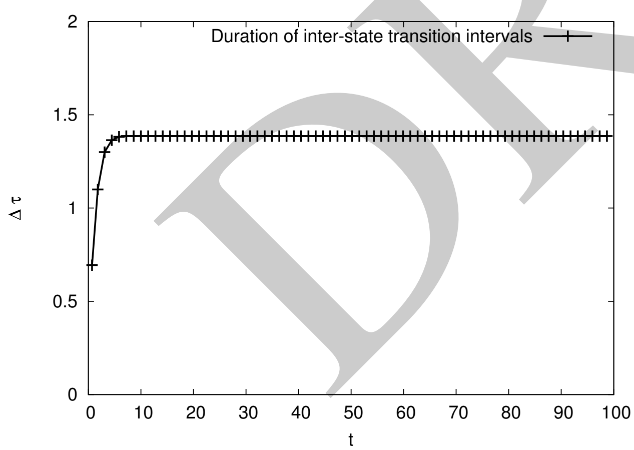

Figure 5: Duration of the inter-state transition intervals over time for a system of 10 units initialized with $S_i(0) = 1/2$ for $i = 1, \ldots, 10$ and $X_i(0) = 0$, for $1 \leq i \leq 5$, and $X_j(0) = 1$ for $6 \leq j \leq 10$. Additionally, $\alpha = 1$, $\mu = 0.25$, $\lambda = 0.75$

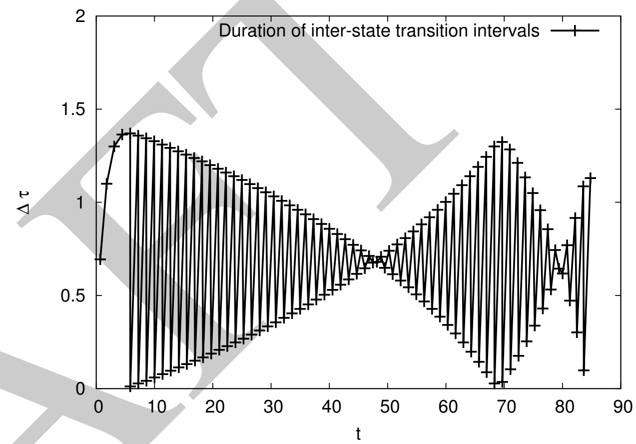

Figure 6: Duration of the inter-state transition intervals over time for a system of 10 units initialized with $S_i(0) = 1/2$ for $i = 1, \ldots, 10$ and $X_i(0) = 0$, for $1 \leq i \leq 3$, and $X_j(0) = 1$ for $4 \leq j \leq 10$. Additionally, $\alpha = 1$, $\mu = 0.25$, $\lambda = 0.75$

Example 2: 0111 (asymmetric blocks of 0's and 1's)

Now consider an initial pattern `0111`, that is, $X_1(0) = 0$, $X_2(0) = 1$, $X_3(0) = 1$, and $X_4(0) = 1$. Also, assume that $S_i(0) = 1/2$ for $i \in \{1, 2, 3, 4\}$. Table 3 shows the evolution of the system for the first 4 transitions.

This example shows that when there are blocks of 0's or 1's of asymmetric length, there are eventually asymmetric initial conditions on the agents' tendencies that makes one agent (in our example, agent 3) switch state before another agent (in our case, agent 4). Through this process, the state of blocks of longer length propagate through the network, which eventually leads to consensus on the initial majority.

The duration of the inter-state transition intervals in the case of asymmetric blocks of 0's or 1's is no longer simple as in the case of symmetric blocks as shown in Figure 6. This figure shows how state updates can become out of sync in asymmetric settings.

It should be clear that in our proposed protocol, informa-

Table 3: Punctuated evolution of a system of 4 units initialized with $S_i(0) = 1/2$ for $i \in \{1, 2, 3, 4\}$ and $X_1(0) = 0$, $X_2(0) = 1$, $X_3(0) = 1$, and $X_4(0) = 1$.

t	X_1	S_1	X_2	S_2	X_3	S_3	X_4	S_4
0	0	$\frac{1}{2}$	1	$\frac{1}{2}$	1	$\frac{1}{2}$	1	$\frac{1}{2}$
$\tau_1 = \frac{1}{\alpha}\ln(\frac{1}{2\mu})$	1	λ	0	μ	1	λ	1	λ
$\tau_2 = \tau_1 + \frac{1}{\alpha}\ln(\frac{\lambda}{\mu})$	1	$1-\frac{\mu^2}{\lambda}$	1	λ	0	μ	1	$1-\frac{\mu^2}{\lambda}$
$\tau_3 = \tau_2 + \min\{\frac{1}{\alpha}\ln(\frac{\lambda}{\mu}), \frac{1}{\alpha}\ln\left(\frac{1-\frac{\mu^2}{\lambda}}{\mu}\right)\} = \tau_2 + \frac{1}{\alpha}\ln(\frac{\lambda}{\mu})$	1	$1-\frac{\mu^3}{\lambda^2}$	1	$1-\frac{\mu^2}{\lambda}$	1	λ	1	$\frac{\mu}{\lambda}\left(1-\frac{\mu^2}{\lambda}\right)$
∞	1	1	1	1	1	1	1	1

tion propagates from one unit to another as waves that either keep their period, or that changes over time depending on whether there are equal numbers of 0's or 1's in the system.

Let us now address the question of whether it is possible for a localized region of concensus to exist or not. That is, is it possible to sustain a configuration where $X_i = 0$ for all i except for a finite set of contiguous agents? We assume for the analysis that follows that each agent is influenced by the next consecutive agent: $j = i + 1$ for all j.

To simplify the problem, we will search for propagating solutions where initially $X_i = 1$ for $i \in [1, m]$. We seek solutions for which the set propagates synchronously to the left, i.e., X_0 transitions to 1 as X_m transitions to 0, and determine what boundary conditions S_0 at $i = 0$ and S_m at $i = m$ are required for a sustained synchronized translation to the left as shown in Figure 7.

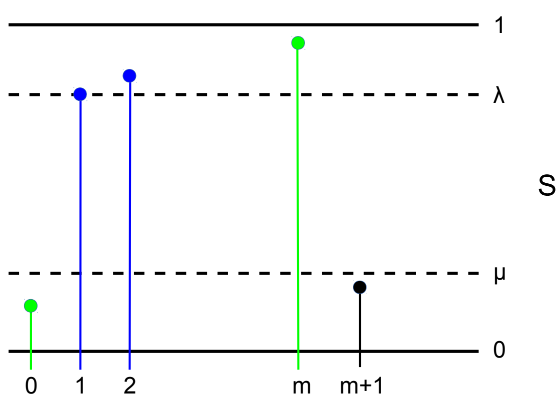

Figure 7: A schematic diagram of a localized region of consensus. The values of S for the translating region of 1's are shown in blue. The boundary agents are indicated in green.

At the leading edge of the region of consensus, we can determine the time Δt_l required for X_0 to transition to 1 after agent 1 transitions to a 1 ($S_1 = \lambda$). In fact, we immediately see that for a translating region of consensus,

$$S_k = 1 - \mu e^{-\alpha(k-1)\Delta t_l}, \quad k \in [1, m]. \tag{12}$$

Also, we see the transition must satisfy the relation,

$$\lambda = 1 - (1 - S_0)e^{-\alpha \Delta t_l} \tag{13}$$

because $X_1 = 1$. We note that S_0 is an imposed boundary condition. Therefore,

$$\Delta t_l = \frac{1}{\alpha}\ln\left(\frac{1 - S_0}{\mu}\right). \tag{14}$$

Meeting the requirement that $\Delta t_l > 0$ is nothing more than requiring $S_0 < \lambda$ which is self-consistent with a transition from $X_0 = 0$ to $X_0 = 1$.

At the trailing edge of the region of consensus, we perform the same calculation. In this case, we can calculate the time Δt_t required for X_m to transition to 0 from 1. Similarly, this event must satisfy the relation

$$\mu = S_m e^{-\alpha \Delta t_t}, \tag{15}$$

so that

$$\Delta t_t = \frac{1}{\alpha}\ln\left(\frac{S_m}{\mu}\right). \tag{16}$$

To enforce a synchronous transition, we require $\Delta t_l = \Delta t_t = \Delta t$, so that the boundary conditions satisfy $S_0 = 1 - S_m$:

$$S_0 = \mu e^{-\alpha(m-1)\Delta t}. \tag{17}$$

Substituting (14) into (17) and simplifying, we can determine S_0 implicitly:

$$\frac{S_0}{(1 - S_0)^{m-1}} = \mu. \tag{18}$$

The solution described by (12), (14) with (18) is artificial in the sense that we cannot specify values of S for the leading agents and regardless they are evolving in time. However, if we consider a cyclic network geometry of length $2m$, we can construct a translating region using this solution for elements 1 through m and its image with $X_k = 0$ and $S_k = 1 - S_{k-m}$ for $k \in [m + 1, 2m]$. In fact, one can construct arrays of such solutions on networks of length $2m$, $4m$ and so forth. The structure of these solutions and lack of free parameters suggest that localized consensus is only possible when there are equal numbers of 0's and 1's.

As we noted earlier, slight perturbations of these solutions lose their synchronization, so the leading and trailing edges will not transition at the same time, and many open questions remain about the stability and evolution of these configurations.

Conclusions

The consensus protocol presented in this paper is derived from previous work of ours on collective decision-making in swarms (Montes de Oca et al., 2012). While the protocol may be seen as a simplification of the full model presented in that work, the protocol exhibits a surprisingly rich dynamical behavior, expecially when the number of units with initial state equal to 0 and 1 are not the same. In this case, the protocol effectively solves the density classification problem. We also showed that when the number of initial 0's and 1's is the same, distributed either as large contiguous blocks of 0's or 1's or as interspersed blocks of equal length, the information waves generated by the protocol travel indefinitely across the network of units that form the system.

We are currently working on analytical solutions that could help us answer questions regarding the system's ability to classify strings of any length, and with arbitrary distributions of 0's and 1's within these strings. Also of interest is the determination of the time needed for convergence.

References

Alonso-Sanz, R. and Bull, L. (2009). A very effective density classifier two-dimensional cellular automaton with memory. *Journal of Physics A: Mathematical and Theoretical*, 42(48):485101.

Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. Santa Fe Institute Studies on the Sciences of Complexity. Oxford University Press, New York.

Deneubourg, J.-L., Goss, S., Franks, N. R., Sendova-Franks, A., Detrain, C., and Chrétien, L. (1990). The dynamics of collective sorting robot-like ants and ant-like robots. In *Proceedings of the first international conference on simulation of adaptive behavior. From animals to animats*, pages 356–363. MIT Press, Cambridge, MA.

Dorigo, M. and Birattari, M. (2007). Swarm Intelligence. *Scholarpedia*, 2(9):1462.

Fatès, N. (2013). Stochastic Cellular Automata Solutions to the Density Classification Problem. When Randomness Helps Computing. *Theory of Computer Systems*, 53(2):223–242.

Fukś, H. (1997). Solution of the density classification problem with two cellular automata rules. *Physical Review E*, 55(3):R2081–R2084.

Fukś, H. (2002). Nondeterministic density classification with diffusive probabilistic cellular automata. *Physical Review E*, 66(6):066106.

Gardner Jr., E. S. (2006). Exponential smoothing: The state of the art–Part II. *International Journal of Forecasting*, 22(4):637–666.

Goss, S., Aron, S., Deneubourg, J.-L., and Pasteels, J.-M. (1989). Self-organized Shortcuts in the Argentine Ant. *Naturwissenschaften*, 76(12):579–581.

Grassé, P.-P. (1959). La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes sp*. La théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes Sociaux*, 6(1):41–80.

Land, M. and Belew, R. K. (1995). No Perfect Two-State Cellular Automata for Density Classification Exists. *Physical Review Letters*, 74(25):5148–5150.

Mesbahi, M. and Egerstedt, M. (2010). *Graph Theoretic Methods in Multiagent Networks*. Princeton Series in Applied Mathematics. Princeton University Press, Princeton, NJ.

Mitchell, M., Crutchfield, J. P., and Hrabel, P. T. (1994). Evolving Cellular Automata to Perform Computations: Mechanisms and Impediments. *Physica D*, 75:361–391.

Montes de Oca, M. A., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., and Dorigo, M. (2011). Majority-Rule Opinion Dynamics with Differential Latency: A Mechanism for Self-Organized Collective Decision-Making. *Swarm Intelligence*, 5(3-4):305–327.

Montes de Oca, M. A., Ferrante, E., Scheidler, A., and Rossi, L. F. (2012). Binary Consensus via Exponential Smoothing. In Glass, K. et al., editors, *Proceedings of the 2nd International Conference on Complex Sciences: Theory and Applications (COMPLEX'12)*, pages 244–255, Berlin, Germany. Springer.

Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., and Laland, K. N. (2010). Why Copy Others? Insights from the Social Learning Strategies Tournament. *Science*, 328(5975):208 – 213.

Visintin, A. (2006). Mathematical Models of Hysteresis. In Bertotti, G. and Mayergoyz, I. D., editors, *The Science of Hysteresis*, volume I, chapter 1, pages 1–114. Academic Press, Oxford, UK.