EUSIPCO 2013 1569747253

REACHING BAYESIAN CONSENSUS IN COGNITIVE SYSTEMS BY DECISION EXCHANGES

Yunlong Wang and Petar M.Djurić

Department of Electrical and Computer Engineering
Stony Brook University, Stony Brook, NY 11794
Email: {yunlong.wang, petar.djuric}@stonybrook.edu

ABSTRACT

We consider the problem of distributed Bayesian hypothesis testing on a set of time invariant hypotheses in a cognitive system of cooperative agents. Each agent in the system obtains one set of private observations and then at every time slot two randomly selected agents repeatedly exchange their decisions and update their beliefs. We propose a method that allows the agents to reach an optimal local consensus by just exchanging decisions. It can be shown that with this strategy, all agents in the system can achieve a consensus in decision, which is also the global optimal decision held by a fictitious fusion center. We provide performance and convergence analysis of the proposed method as well as simulation results that demonstrate its asymptotical properties.

***Index Terms*—** Cognitive system, Bayesian consensus, distributed hypothesis testing, cooperative agents

1. INTRODUCTION

In this paper we address the problem of reaching a Bayesian consensus in a cognitive system modeled by a set of spatially distributed agents. The agents are linked through a connected graph where any two neighboring agents can exchange their decisions without any noise in communication. The agents receive private signals generated according to one of two or more possible models (hypotheses), all known to the agents. After receiving the private observations, all agents repeatedly make decisions about the hypotheses, learn from the decisions of their neighbors, and adapt their beliefs until all agents reach consensus identical to the optimal global decision.

In [1], it was maintained that an agent with cognitive capacity should be able to learn from the environment, including the information from other agents, and adapt its internal states. An important question was raised: in what way should the agent learn from environmental information and in what way should its inner state be adapted? This has been widely studied in the realm of cognitive radio [2, 3] and cognitive control [4]. In the studies of cognitive systems where cognition is achieved in a distributed manner, the learning and adaptation processes are cooperative. For example, in [5] and [6], the authors propose cooperative spectrum sensing methods in cognitive radio.

In this work, we address the problem of Bayesian learning in cooperative networks [7, 8], where the agents in the network asymptotically attain the performance of a Bayesian fusion center by distributed cooperation. We focus on applications where the information exchanged between agents are just their decisions on the true state of nature. In [9], we studied the case of binary hypotheses, whereas in this paper, we extend our work to multiple hypotheses. More specifically, we propose a solution based on the idea of making the beliefs of all the agents converge to an identical decision region instead of to an identical value. We analyze the performance of the method theoretically and demonstrate it by simulations.

The paper is organized as follows. In the next section we state the problem. In Section 3, we present the solution when the agents choose from two hypotheses, and in Section 4, from multiple hypotheses. Section 5 provides simulation results, and Section 6 contains conclusions.

2. PROBLEM STATEMENT

Mathematically, we formulate the problem as follows. We consider a distributed hypothesis testing problem in a cognitive system which includes N cooperative agents A_i, $i \in \mathcal{N}_A = \{1, 2, ..., N\}$. The connections among agents are described by an undirected graph $G = (\mathcal{N}_A, \mathcal{E})$, where $\mathcal{E}$ is the set of edges of the graph, and where agents A_i and A_j can directly exchange information if and only if $(i, j) \in \mathcal{E}$. In this paper, we assume that the topology of the network is time invariant and that the communication between any two communicating agents is perfect.

In this cognitive system, every agent A_n has its private observations z_n that may be generated according to one of K hypotheses, which are denoted by $\{\mathcal{H}_1, \mathcal{H}_2, \cdots, \mathcal{H}_K\}$. We assume that the private observations are generated from the same hypothesis and that the observations of different agents are independent. At time slot $t = 1$, for all $n \in \mathcal{N}_A$, agent A_n can obtain its beliefs on each hypothesis $\mathcal{H}_k$, $k \in \mathcal{N}_H =$

This work was supported by NSF under Award CCF-1018323.

$\{1, 2, ..., K\}$ by using Bayes' rule, i.e., it can form the vector of posterior probabilities

$$\boldsymbol{\pi}_n[1] = [\pi_{n,1}[1], \pi_{n,2}[1], ..., \pi_{n,K}[1]]^\top, \quad (1)$$

where $\pi_{n,k}[1] = P(\mathcal{H}_k|z_n)$. The agents are allowed to repeatedly make decisions on the state of nature and modify their private beliefs by using their neighbors' decisions. At time slot $t \in \mathbb{N}^+$, for any $n \in \mathcal{N}_A$, the A_n's belief is denoted by the vector $\boldsymbol{\pi}_n[t] = [\pi_{n,1}[t], \pi_{n,2}[t], ..., \pi_{n,K}[t]]^\top$, where $\pi_{n,k}[t]$ is the A_n's belief in $\mathcal{H}_k$ at time slot t.

Let the reward of any agent A_n, $n \in \mathcal{N}_A$ be `one` if its decision is identical to the true hypothesis, and `zero` otherwise; then the A_n's expected reward at t is maximized if it makes decision $\alpha_n[t]$ by the following rule:

$$\alpha_n[t] = \arg_k \max \pi_{n,k}[t]. \quad (2)$$

We also consider a fictitious fusion center that has access to the initial beliefs of all the agents in the network, $\pi_{n,k}[1]$, $\forall n \in \mathcal{N}_A$ and $\forall k \in \mathcal{N}_H$. These beliefs can be fused using Bayes' theorem to form an *optimal* belief $\boldsymbol{\pi}_o = [\pi_{o,1}, \pi_{o,2}, ..., \pi_{o,K}]^\top$, where for each $k \in \mathcal{N}_H$, $\pi_{o,k}$ denotes the posterior probability of $\mathcal{H}_k$. Each element of the *optimal* belief is given by [7]

$$\pi_{o,k} \propto \prod_{n=1}^{N} \pi_{n,k}[1]. \quad (3)$$

We assume that the reward of the fictitious fusion center, too, is `one` if it chooses the correct hypothesis, and `zero`, otherwise. Thus, its expected reward is also maximized if $\alpha_o = k$ given that $\pi_{o,k}$ is greater than any other element in $\boldsymbol{\pi}_o$. Here we assume that $\forall n \in \mathcal{N}_A$ and $\forall k \in \mathcal{N}_H$, the beliefs $\pi_{n,k}[1]$ are continuous random variables, and therefore the probability that the optimal decision is not unique is zero. Therefore, in the analysis, we address the case of unique optimal decisions only. We use the performance of the fictitious fusion center as a benchmark for the studied system.

According to [7], the result in (3) implies that every A_i's decision will be identical to the optimal decision of the fictitious fusion center if A_i can obtain the average value of all the agents' log-beliefs on every hypothesis $\mathcal{H}_k$, $k \in \mathcal{N}_H$,

$$\bar{\boldsymbol{l}} = \left[\frac{1}{N}\Sigma_{n=1}^{N} l_{n,1}[1], \frac{1}{N}\Sigma_{n=1}^{N} l_{n,2}[1], \cdots, \frac{1}{N}\Sigma_{n=1}^{N} l_{n,K}[1]\right]^\top,$$

where $l_{n,i}[1]$ is the log-belief of agent A_i at $t = 1$.

The problem of achieving average value in distributed manner can be solved by the average consensus [10] or gossip algorithms [11] if the agents are allowed to exchange their log-beliefs. In this paper, however, we study the problem where the communication between the agents is constrained to only a few quantized values. We note that neither a quantized average consensus algorithm nor a quantized gossip algorithm [12] can guarantee the consensus in decision.

In the following sections, we propose an algorithm by which all the agents' decisions converge to the optimal one held by a fictitious fusion center and where the agents repeatedly exchange only their decisions, i.e., an algorithm with the following property:

$$\lim_{t\to\infty} P(\alpha_n[t] = \alpha_o) = 1, \forall\, n \in \mathcal{N}_A. \quad (4)$$

Before we proceed, we reiterate that our system is cognitive and is composed of cooperative agents. The behavior of the agents is modeled by a closed cognitive cycle, where at each time slot only two agents implement cooperation, as in a gossip style setting. Namely, we assume that every agent A_i is waken up with probability $\frac{1}{N}$, and then the agent selects one of its neighbors, e.g., A_j with a positive probability $P_{i,j}$ to exchange decisions. With the exchanges, these agents update their log-beliefs until they reach a local consensus in their decisions. In Fig. 1, we describe pictorially the behavior of every agent in this cognitive system.

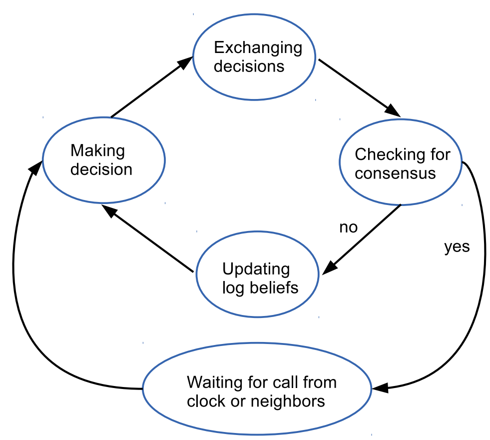

Fig. 1. The model of agent behavior in the cognitive system.

3. THE BINARY HYPOTHESIS CASE

In this section, we introduce the proposed method on the binary hypothesis testing problem, where $\mathcal{N}_H = \{1, 2\}$. We propose a gossip-type algorithm where the agents exchange merely their decisions and yet, the probability that every agent reaches the optimal consensus in decision converges to one.

Let the log-belief ratio (LBR) of agent A_n at time slot t be $\beta_n[t] = l_{n,2}[t] - l_{n,1}[t]$. The decision is made based on the LBR of the agent, that is, if the LBR of A_n is above a threshold or below it. The idea behind the algorithm is to get a global optimal consensus by reaching optimal local consensuses at each time slot. With the same settings as for the gossip algorithm, at each time slot, a pair of neighboring agents A_i and A_j is selected with probability $(P_{ij} + P_{ji})/N$. The two selected agents, implement an algorithm referred to as the Local Consensus Algorithm (LCA), which allows them to reach an optimal local consensus after a finite number of exchanges of their decisions [9].

We define the variables $U_n^{(m)}$ and $L_n^{(m)}$, which represent the upper and lower bounds of the LBR $\beta_n[t]$ of A_n,($n \in \{i,j\}$), at iteration m, respectively.[1] We also define the threshold $\gamma_n^{(m)}$ that is used for decision making at iteration m (and time slot t). During the implementation of the LCA, the bounds $L_n^{(m)}$ and $U_n^{(m)}$ and the threshold $\gamma_n^{(m)}$ are all computed by A_n. We formally describe the LCA as follows:

Initialization: For A_n, $n \in \{i,j\}$, set $\gamma_n^{(1)} = 0$, $U_n^{(0)} = \infty$, $L_n^{(0)} = -\infty$ and $\beta_n[t] = l_{n,2}[t] - l_{n,1}[t]$. The following steps describe the mth iteration.

Step 1 (Intermediate decision exchanges) The agents make intermediate decisions, $\delta_n^{(m)}$, $n \in \{i,j\}$, based on their individual log-beliefs and thresholds, i.e.,

$$\delta_n^{(m)} = \begin{cases} 2, & \text{if } \beta_n[t] > \gamma_n^{(m)} \\ 1, & \text{if } \beta_n[t] < \gamma_n^{(m)} \end{cases}. \tag{5}$$

Then A_i transmits its intermediate decision $\delta_i^{(m)}$ to A_j, and A_j sends its $\delta_j^{(m)}$ to A_i.

Step 2: (Consensus checking) If the agents achieve consensus, i.e., $\delta_i^{(m)} = \delta_j^{(m)}$, they stop, and they set their log-beliefs to

$$l_{n,1}[t+1] = l_{n,1}[t] + \frac{\gamma_n^{(m)}}{2} \tag{6}$$

$$l_{n,2}[t+1] = l_{n,2}[t] - \frac{\gamma_n^{(m)}}{2}, \tag{7}$$

and by (2), the agents have a local consensus $\alpha_n[t+1] = \delta_n^{(m)}$. Otherwise, they proceed to the next step.

Step 3: (Interval updating) The agents update their intervals L_n and U_n according to

$$\left(L_n^{(m+1)}, U_n^{(m+1)}\right) = \begin{cases} \left(\gamma_n^{(m)}, U_n^{(m)}\right), & \text{if } \delta_n^{(m)} = 2 \\ \left(L_n^{(m)}, \gamma_n^{(m)}\right), & \text{if } \delta_n^{(m)} = 1 \end{cases}. \tag{8}$$

Step 4: (Threshold updating) Each agent updates its own threshold by $\gamma_n^{(m+1)} = f\left(L_n^{(m)},\ U_n^{(m)}\right)$, where the function $f(\cdot,\cdot)$ is identical and known to every agent.

Before proceeding with the analysis of the algorithm, we provide an example.

[1] Note that in order to avoid a too unwieldy notation, we do not specify that the bounds are for time slot t. A full notation would be, e.g., $U_n^{(m)}[t]$.

Example 1 *Let A_i and A_j be selected at time slot t, and let the rule for obtaining the threshold $\gamma_n^{(m)}$ be defined by*

$$\gamma_n^{(m+1)} = \begin{cases} \gamma_n^{(m)} + \left(\delta_n^{(m)} - 3/2\right)\Delta, & \text{if } \mathcal{C} \text{ is true} \\ \dfrac{L_n^{(m)} + U_n^{(m)}}{2}, & \text{otherwise,} \end{cases} \tag{9}$$

where $\mathcal{C}$ stands for the condition $L_n^{(m)} = -\infty$ or $U_n^{(m)} = \infty$, and Δ is a positive real number. Suppose also that $\beta_i[t] = 0.412\Delta$ and $\beta_j[t] = -0.207\Delta$. Therefore $\beta_i[t] + \beta_j[t] > 0$, and the optimal local consensus is $\alpha_o[t] = 2$. We refer the reader to a similar example in [9] where we show in more detail how the agents reach the optimal decision.

In [9], we have shown that under mild conditions, by using the LCA algorithm all the agents asymptotically reach the optimal decision of a fictitious fusion center. We formally state this analytical result with the following theorem:

Theorem 1 *Let the initial log-beliefs $\mathbf{l}_n[1] \in \mathbb{R}^{2\times 1}$, $\forall n \in \mathcal{N}_A$, be continuous random vectors. With the LCA-based gossip algorithm the agents achieve the optimal consensus in probability, i.e., $\lim_{t\to\infty} P(\boldsymbol{\alpha}[t] = \alpha_o\ \mathbf{1}) = 1$, where $\mathbf{1} \in \mathbb{R}^{N\times 1}$ is a vector of all ones, for as long as the threshold update function of the LCA $f(\cdot,\cdot)$ (Step 4 of the LCA) has the following two properties:*

(P1) $\gamma_i^{(m)} = -\gamma_j^{(m)}$, for any $m = 1, 2, 3, \cdots$,

(P2) *If $\delta_i^{(m)} \neq \delta_j^{(m)}$ for any $m \in \{1, 2, 3, \cdots, M\}$, then* $\lim_{M\to\infty} U_n^{(M)} - L_n^{(M)} = 0$, $\forall n \in \{i,j\}$.

Furthermore, for as long as (P1) and (P2) hold, the LCA can preserve the mean value of the log-beliefs and guarantee a local optimal consensus in a finite number of iterations.

Theorem 2 *In the execution of the LCA, if* (P1) *and* (P2) *hold, $\forall k \in \mathcal{N}_H$, the sum of all the log-beliefs on $\mathcal{H}_k$ in the network does not change with t, i.e.,*

$$\sum_{n=1}^{N} l_{n,k}[t+1] = \sum_{n=1}^{N} l_{n,k}[t],\ \forall k \in \{1,2\}, \tag{10}$$

and there exists a finite iteration $\tilde{m}$ such that $\alpha_n[t+1] = \delta_i^{(\tilde{m})} = \delta_j^{(\tilde{m})} = \delta_o$, $\forall n \in \{i,j\}$, where

$$\delta_o = \begin{cases} 2, & \text{if } \dfrac{\beta_i[t] + \beta_j[t]}{2} > 0 \\ 1, & \text{if } \dfrac{\beta_i[t] + \beta_j[t]}{2} < 0 \end{cases}. \tag{11}$$

According to Theorem 1, since every agent can eliminate one hypothesis by implementing the LCA, it can be expected that when $K > 2$, all agents will also reach the optimal consensus by repeating the LCA $K-1$ times. However, in the following section, we propose a method achieving the same goal with much less communication.

4. THE MULTIPLE HYPOTHESIS CASE

Now, we consider the problem when the number of hypotheses K is greater than two. We propose a gossip-type algorithm where the agents repeatedly exchange decisions about their favoring hypotheses. In other words, an agents sends to its communicating agent the index of the hypothesis with the largest log-belief. Furthermore, each agent adapts its beliefs in the two hypotheses chosen by these agents.

Let A_i and A_j be selected at time slot t. They carry out an extended local consensus algorithm (ELCA), where the agents repeatedly implement LCA until they reach the local optimal consensus. In each iteration, the agents make their temporary decision[2] and check whether the local consensus has been reached. If consensus has not been attained, they implement the LCA to adapt their temporary log-beliefs and find out which of the two decisions is better. We also denote the temporary log-belief by $l_{n,k}^{(m)}$, which represents the temporary log-belief on the kth hypothesis held by agent A_n at iteration m at time slot t, and we set $l_{n,k}^{(1)} = l_{n,k}[t]$. At the mth iteration, the agent A_n, $\forall n \in \{i, j\}$ carries out the following steps:

Step 1 (Temporary decision making): $\forall n \in \{i, j\}$, agent A_n makes its mth temporary decision by

$$\alpha_n^{(m)} = \arg_k \max l_{n,k}^{(m)}, \quad \forall k \in \{1, 2, ..., K\}.$$

Agent A_i then transmits its temporary decision $\alpha_i^{(m)}$ to A_j, and A_j sends its temporary decision $\alpha_j^{(m)}$ to A_i.

Step 2 (Consensus checking): If $\alpha_i^{(m)} = \alpha_j^{(m)}$, they stop the communication and set $l_{n,k}[t+1] = l_{n,k}^{(m)}$, $\forall k \in \mathcal{N}_H$, and by (2) they have their decisions $\alpha_i[t+1] = \alpha_i^{(m)}$ and $\alpha_j[t+1] = \alpha_j^{(m)}$. Otherwise, they proceed to the next step.

Step 3 (LCA implementation): Let $\alpha_b = \min(\alpha_i^{(m)}, \alpha_j^{(m)})$ and $\alpha_B = \max(\alpha_i^{(m)}, \alpha_j^{(m)})$. The agent A_n then adapts its temporary log-belief from $l_{n,\alpha_B}^{(m)}$ and $l_{n,\alpha_b}^{(m)}$ to $l_{n,\alpha_B}^{(m+1)}$ and $l_{n,\alpha_b}^{(m+1)}$ by implementing the LCA with input LBR $l_{n,\alpha_B}^{(m)} - l_{n,\alpha_b}^{(m)}$.

Again, we use an example to illustrate the method.

Example 2 *Let A_i and A_j be selected at time slot t and and let their log-beliefs be $\mathbf{l}_i[t] = [-0.9, -0.2, -0.3]^\top$ and $\mathbf{l}_j[t] = [-0.3, -0.7, -0.4]^\top$. According to (2), it is shown that the optimal local consensus should be $\mathcal{H}_3$ because the mean value of the log-beliefs of $\mathcal{H}_3$ is the largest.*

In the first iteration, by step 1, A_i and A_j exchange their first temporary decisions $\alpha_i^{(1)} = 2$ and $\alpha_j^{(1)} = 1$. After step 2, they find that the local consensus has not been reached and so they proceed to step 3. At step 3, both agents have $\alpha_b = 1$ and $\alpha_B = 2$ and they implement the LCA with $l_{i,2}[t] - l_{i,1}[t] = 0.7$ and $l_{j,2}[t] - l_{j,1}[t] = -0.4$. Let the threshold update rule of the LCA be given by the update rule in Example 1, and let Δ be 1. Then we get that the log-beliefs after the adaptation in step 3 are

$$l_{i,2}^{(2)} = l_{i,2}^{(1)} - 0.25 = -0.45, \tag{12}$$

$$l_{j,2}^{(2)} = l_{j,2}^{(1)} + 0.25 = -0.45, \tag{13}$$

$$l_{i,1}^{(2)} = l_{i,1}^{(1)} + 0.25 = -0.65, \tag{14}$$

$$l_{j,1}^{(2)} = l_{j,1}^{(1)} - 0.25 = -0.55. \tag{15}$$

In the second iteration, the log-beliefs have been updated to $\mathbf{l}_i^{(2)} = [-0.65, -0.45, -0.3]^\top$ and $\mathbf{l}_j^{(2)} = [-0.55, -0.45, -0.4]^\top$, and the temporary decisions of the agents are $\alpha_i^{(2)} = 3$, and $\alpha_j^{(2)} = 3$. At step 2, the agents determine that they have reached a local consensus, and thus they update their log-beliefs to $\mathbf{l}_i[t+1] = [-0.65, -0.45, -0.3]^\top$ and $\mathbf{l}_j[t+1] = [-0.55, -0.45, -0.4]^\top$ and set their decisions to $\alpha_i[t+1] = \alpha_j[t+1] = 3$.

We note that if the two conditions in Theorem 1 hold, according to Theorem 2, the average log-belief of the two agents in each hypothesis is constant. The latter guarantees that once a local consensus is reached, it must be the optimal consensus between the two agents. Also, since in every iteration of the ELCA the two agents can eliminate one hypothesis by implementing the LCA, it takes the agents at most $K-1$ iterations to reach a local consensus. Finally, as shown in the simulations, the ELCA-based method has a finite consensus time T_c such that $\forall t > T_c$ and $\forall n \in \mathcal{N}_A$, the decision of all the agents is $\alpha_n[t] = \alpha_o$, which is the decision of the fictitious fusion center.

5. SIMULATION

In this section, we provide computer simulations that show the convergence performance of the ELCA-based method. The multi-agent cognitive system was modeled as a random geometric graph $G(\mathcal{N}_A, \mathcal{E})$, where the N agents were chosen uniformly and independently on a square of size 1×1. Each pair of agents was set to be connected if the Euclidean distance between them was smaller than $r(N) = \sqrt{\frac{\log(N)}{N}}$ due to connectivity requirement. In the experiment, the connectivity was checked.

We considered the following hypothesis test: the agents observed data and they had to choose between $K = 10$ hypotheses. We assumed that the agents observed data from one of the following Gaussian hypotheses:

$$\mathcal{H}_k : \quad z_n \sim N(\theta_k, \sigma_w^2), \ \forall k \in \mathcal{N}_H = \{1, \cdots, 10\}, \tag{16}$$

[2] They are called temporary because they are used for consensus checking.

where $\theta_k = k\theta$ was known and the random perturbation was modeled as a zero mean Gaussian random variable with known variance σ_w^2. For the prior probabilities of the hypotheses, we let $P(\mathcal{H}_k) = 1/10$. Without loss of generality, we assumed that the data were generated from $\mathcal{H}_1$, and we set $\theta = 1$, and $\sigma_w = 5$. Note that (16) is a very simple model; however, the ELCA-based method can readily be used for any model where the agents can obtain their initial log-beliefs.

In the experiment, we had $N = 20$ agents, which observed data generated from (16) and tried to reach consensus by the proposed method. In implementing the ELCA, the agents used the update rule from Example 1 with different values of Δ, where $\Delta_1 = 0.5$, $\Delta_2 = 2$, $\Delta_3 = 4$, respectively. We ran the experiment 1000 times.

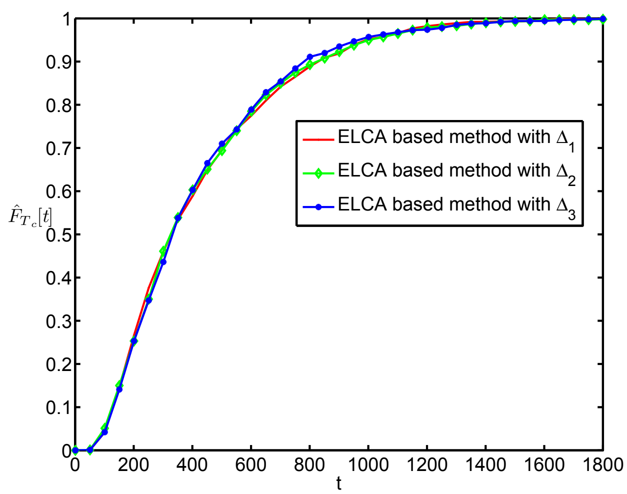

Fig. 2. Cumulative distribution functions of consensus times for ELCA with different parameters.

At each time slot t, we defined $C[t]$ as the number of trials where the agents reached the optimal consensus before t. Then $C[t]/1000$ is an approximation of the cumulative distribution function of the consensus time T_c, i.e, $F_{T_c}[t] = P(T_c \leq t)$. The results are shown in Fig. 2, where we see that the system using ELCA-based methods had finite T_C in all the trials. It can also be seen that for different values of Δ, ELCA has the same asymptotical performance in this simulation. Moreover, we emphasize that in each trial of this experiment, once the consensus is reached among all agents, it is identical with the decision held by the fictitious fusion center.

In this experiment we also studied the communication cost for reaching the local consensus. If we define the cost as the number of decision transmissions by ELCA, then for Δ_1, Δ_2, and Δ_3, the average communication cost was 9.22, 11.31, 12.49, respectively.

6. CONCLUSION

In this paper, we studied a system with Bayesian agents that process data and use them to make decisions on the state of nature. They exchange these decisions in a gossip manner with neighbors until they reach local optimum consensus. When the state of nature is binary, we have a proof that all the agents attain the optimal consensus with probability one. We generalized the method for multiple hypotheses, and showed its performance by Monte Carlo simulations.

7. REFERENCES

[1] S. Haykin et al., "Cognitive radio: Brain-empowered wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 2, pp. 201–220, 2005.

[2] T. Yucek and H. Arslan, "A survey of spectrum sensing algorithms for cognitive radio applications," *Communications Surveys & Tutorials*, vol. 11, no. 1, pp. 116–130, 2009.

[3] A. Jovičić and P. Viswanath, "Cognitive radio: An information-theoretic perspective," *IEEE Transactions on Information Theory*, vol. 55, no. 9, pp. 3945–3958, 2009.

[4] K. N. Ochsner and J. J. Gross, "The cognitive control of emotion," *Trends in Cognitive Sciences*, vol. 9, no. 5, pp. 242–249, 2005.

[5] S. M. Mishra, A. Sahai, and R. W. Brodersen, "Cooperative sensing among cognitive radios," in *Proceedings of the IEEE International Conference on Communications*, 2006, vol. 4, pp. 1658–1663.

[6] M. Gandetto and C. Regazzoni, "Spectrum sensing: A distributed approach for cognitive terminals," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 3, pp. 546–557, 2007.

[7] P. M. Djurić and Y. Wang, "Distributed Bayesian learning in multiagent systems: Improving our understanding of its capabilities and limitations," *IEEE Signal Processing Magazine*, vol. 29, no. 2, pp. 65–76, 2012.

[8] D. Acemoglu, M. A Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[9] Y. Wang and P. M. Djurić, "A gossip method for optimal consensus on a binary state from binary actions," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 2, pp. 274–283, 2013.

[10] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.

[11] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2508–2530, 2006.

[12] A. Kashyap, T. Başar, and R. Srikant, "Quantized consensus," *Automatica*, vol. 43, no. 7, pp. 1192–1203, 2007.