Reaching Consensus with Imprecise Probabilities over a Network

Luca F. Bertuccelli* and Jonathan P. How†

ACL Report 08-02

Aerospace Controls Laboratory
Massachusetts Institute of Technology

December 20, 2008

Abstract

This paper discusses the problem of a distributed network of agents attempting to agree on an imprecise probability over a network. Unique from other related work however, the agents must reach agreement while accounting for relative uncertainties in their respective probabilities. First, we assume that the agents only seek to agree to a centralized estimate of the probabilities, without accounting for observed transitions. We provide two methods by which such an agreement can occur which uses ideas from Dirichlet distributions. The first methods interprets the consensus problem as an aggregation of Dirichlet distributions of the neighboring agents. The second method uses ideas from Kalman Consensus to approximate this consensus using the mean and the variance of the Dirichlet distributions. A key results of this paper is that we show that when the agents are simultaneously actively observing state transitions and attempting to reach consensus on the probabilities, the agreement protocol can be insensitive to any new information, and agreement is not possible. Ideas from exponential fading are adopted to improve convergence and reach a consistent agreement.

1 Introduction

Many large-scale systems can be treated as a team of heterogeneous agents with local information such as military command and control architectures, unmanned aerial vehicle missions, or interactions among medical groups. Each of the agents can be treated as a node of a more complex network, and due to physical proximity to a set of events, can have localized information that may be unavailable directly to other members of the network. As a result, it is of key importance to share the information across the network, and be able to reach some agreement with respect to the nature of the observed event. In particular, each agent has an imprecise probability they are trying to bring agreement with the rest of the network. The probability can be imprecise for various reasons, one of them being that an insufficient number of measurements has been taken that allow the agent to conclusively estimate the probability.

A particular instance in which this agreement is a requirement is a cooperative classification tasks with a mixed team of humans and UAVs, where a human would desire agreement among multiple agents on the probability that the object being classified is indeed a target. Another instance is the case of multiple doctors that receive independent results of a test on a patient, and must agree on the patient's state in a decentralized manner (i.e., they cannot all meet up in one central place and come to agreement on the patient's health). Yet another problem of extreme relevance is the multi-agent extension of robust planning using Markov Decision Processes (MDPs) for sequential decision-making, where agents may need to perform robust decisions, but must come to some agreement as to the underlying transition probabilities of the dynamic system.

2 Problem Statement and Previous Work

The objective is to achieve a consensus on a probability vector $\bar{p}(k) \in \Re^M$ in a group of N agents (see Figure 1). In other words, with each agent i having a mean belief, $\bar{p}^i(k)$, the objective is to asymptotically reach an agreement such that $\bar{p}^1(k) = \bar{p}^2(k) = \ldots = \bar{p}^M(k)$. The agent connectivity is described by a graph $G(E,V)$, where E denotes the edges of the graph and V denotes the vertices, and the neighbors (or neighboring nodes) of agent i are denoted by N_i.

*Postdoctoral Associate, Department of Aeronautics and Astronautics: lucab@mit.edu
†Professor, Dept of Aeronautics and Astronautics jhow@mit.edu

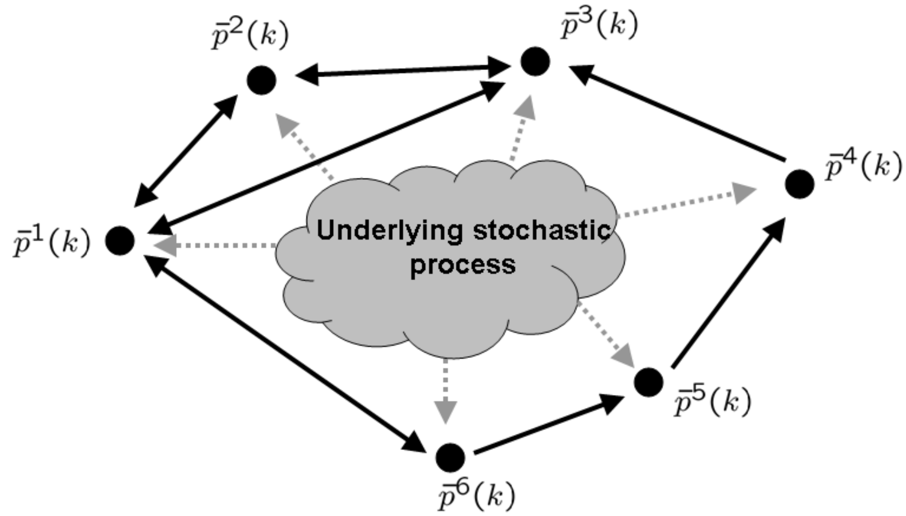

Figure 1: Example of a multi-agent system estimating a probability

Reaching consensus over a network has been investigated extensively in a variety of fields [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]. With regards to agreeing on probabilities, the seminal work of Winkler [9, 10] and DeGroot [3] provides some of the earlier work on agreement on subjective probabilities. However, consensus was not in the context of the more recent literature, in that there was no specific mention of a network structure, nor the limitations of a fully connected implementation. The work by Castañón [1, 2] also discusses agreement on probabilities, but no specific mention of a network is made. Teneketzis and Varaiya [8] and Washburn and Teneketzis [11] also discuss the issue of agreement on the subjective beliefs of distributed decision-makers, but do not discuss the issue of graph connectivity. Genest [5] provides a nice overview of the problem. Belief propagation [12] also is related to this problem, in the context of graphical models.

In a more general consensus framework, where information is provided in a Kalman filtering framework, Ren, Beard and Kingston [13] and Ren, Beard, and Atkins [14] addressed the problem of reaching consensus on information with a measure of uncertainty, and this gave rise to Kalman Consensus algorithms. However, this work did not address the issue of consensus on probabilities. Also, the work of Olfati-Saber et al [15] and Spanos, Olfati-Saber, and Murray [16] discuss the issue of distributed hypothesis testing, and distributed Kalman filtering, respectively but did not address the issue of these beliefs being uncertain or imprecise.

2.1 Consensus

One of the many methods to achieve consensus on an imprecise probability is to use ideas from linear consensus on the probability ([15, 16]), which would achieve consensus only on the mean belief, $\bar{p}$. Then each agent i would run a consensus algorithm on their belief $\bar{p}^i(k)$ of the form:

$$\bar{p}^i(k+1) = \bar{p}^i(k) + \varepsilon \sum_{j \in N_i} (\bar{p}^j(k) - \bar{p}^i(k)) \tag{1}$$

Here N_i is the collection of neighboring nodes of agent i, and ε is a constant (chosen by the designer) that depends on network connectivity. It is shown in [15] that ε is constrained in $\varepsilon \in (0, 1/\Delta(G))$, where $\Delta(G)$ represents the maximal node degree of the network.

A particularly appealing motivation for linear consensus, is that can be shown that a consensus is reached, and, perhaps more importantly, that some key properties of the belief are retained. Namely,

$$\begin{aligned} \mathbf{1}^T \bar{p}^i(k) &= 1, \qquad \forall i,\ k \\ \mathbf{0} \leq \bar{p}^i(k) &\leq \mathbf{1}, \qquad \forall i,\ k \end{aligned} \tag{2}$$

However, such a linear consensus protocol does not incorporate sufficient information on the relative uncertainty of the probabilities. As a simple example, consider Figure 2(a) where two agents (blue and green) both observe realizations of a stochastic process with $M = 3$, such as in a classification task, the probability that the target is a tank, and concluded that the probability of the first event is $p^1 = 0.335$ and $p^2 = 0.27$ respectively. However, the first agent concluded this after having observed 1 out of 3 images, in the classification task, but the second agent concluded this after having observed 27 out of 100 images. Clearly, the linear consensus solution of $p_{LC} = 0.31$ is inconsistent with the underlying observations, and different from the centralized solution that includes the differing number of observations (shown in black).

A formal method needs to be accounted for to embed a more general notion of the uncertainty in the probability, and such an approach is by the use of the Dirichlet distribution.

2.2 Use of the Dirichlet

The Dirichlet distribution [17] f_D on a probability vector $\mathbf{p}_k = [p_1, p_2, \ldots, p_N]^T$ at a time instant k is given by

$$\begin{aligned} f_D(\mathbf{p}_k|\alpha(\mathbf{k})) &= K \prod_{i=1}^{N} p_i^{\alpha_i - 1}, \quad \sum_i p_i = 1 \\ &= K\, p_1^{\alpha_1 - 1} p_2^{\alpha_2 - 1} \ldots (1 - \sum_{i=1}^{N-1} p_i)^{\alpha_N - 1} \end{aligned} \tag{3}$$

where K is a normalizing factor that ensures the probability distribution integrates to unity and a set of hyperparameters (with $\alpha_i > 1$) $\alpha(k) = [\alpha_1, \alpha_2, \ldots, \alpha_N]^T$ define the shape of the distribution.

The primary reasons for using the Dirichlet distribution is that the mean $\bar{p}_i$ satisfies the requirements of a probability vector $0 \leq \bar{p}_i \leq 1$ and satisfies the unit sum constraint $\sum_i \bar{p}_i = 1$ by construction. Furthermore, the hyperparameters α_i can be interpreted as “counts", or the number of times that the particular event was observed.

The mean and the variance of the Dirichlet distribution can be calculated directly as follows

$$\bar{p}_i = \alpha_i / \alpha_0 \tag{4}$$

$$\Sigma_{ii} = \frac{\alpha_i(\alpha_0 - \alpha_i)}{\alpha_0^2(\alpha_0 + 1)} \tag{5}$$

It is well known that the Dirichlet distribution is conjugate to the multinomial distribution [17]; therefore, performing a Bayesian measurement update step on the Dirichlet amounts to a simple addition of currently observed transitions to the previously observed counts $\alpha(k)$. The posterior distribution $f_D(\mathbf{p}_{k+1}|\alpha(k+1))$ is given in terms of the prior $f_D(\mathbf{p}_k|\alpha(k))$ as

$$\begin{aligned} f_D(\mathbf{p}_{k+1}|\alpha(k+1)) &\propto f_D(\mathbf{p}_k|\alpha(k)) f_M(\beta(k)(|\mathbf{p}_k) \\ &= \prod_{i=1}^{N} p_i^{\alpha_i - 1} p_i^{\beta_i} = \prod_{i=1}^{N} p_i^{\alpha_i + \beta_i - 1} \end{aligned}$$

where $f_M(\beta(k)|\mathbf{p}_k)$ is a multinomial distribution with hyperparameters $\beta(k) = [\beta_1, \ldots, \beta_N]$. Each β_i is the total number of transitions observed from state i to a new state i': mathematically $\beta_{i'} = \sum_i \delta_{i,i'}$ and

$$\delta_{i,i'} = \begin{cases} 1 & \text{if } i = i' \\ 0 & \text{else} \end{cases}$$

Upon receipt of the observations $\beta(k)$, the parameters $\alpha(k)$ are thus updated in the following manner

$$\alpha_i(k+1) = \begin{cases} \alpha_i(k) + \delta_{i,i'} & \text{Transition } i \text{ to } i' \\ \alpha_i(k) & \text{Else} \end{cases}$$

The mean and the variance can then be calculated by using Eqs. 4 and 5.

3 Results on the agreement problem without information updates

3.1 Approach #1: Dirichlet aggregation

One interpretation of an agreement problem is to treat each iteration of the consensus as a fictitious information update from the other agents. That is, at each iteration k each agent i passes its latest updated information aggregated from the other agents in its local neighborhood N_i. Thus, the aggregation of agent i's unique prior $f_D(\mathbf{p}_k|\alpha^i(k))$ and the information of each neighboring agent j can be modeled as

$$f_D(\mathbf{p}_{k+1}|\alpha^i(k+1)) \propto \underbrace{f_D(\mathbf{p}_k|\alpha^i(k))}_{\text{Agent prior Dirichlet}} \times \underbrace{\prod_{j\in N_i} f_D(\mathbf{p}_k|\alpha^j(k))}_{\text{Neighborhood Dirichlet}}$$

and a consensus algorithm that seeks to reach agreement on the mean probability (or mean belief) is Algorithm 1:

Algorithm 1 Dirichlet aggregation

Choose a convergence criterion ε
while Convergence criterion not satisfied **do**
Update Dirichlet: $f_D(\mathbf{p}_{k+1}|\alpha^i(k+1)) \propto f_D(\mathbf{p}_k|\alpha^i(k)) \times \prod_{j\in N_i} f_D(\mathbf{p}_k|\alpha^j(k))$

Find the mean belief:

$$\begin{aligned} \bar{p}^i(k+1) &= \int_0^1 \mathbf{p}_{k+1} f_D(\mathbf{p}_{k+1}|\alpha^i(k+1)) d\mathbf{p}_{k+1} \\ &= \frac{\alpha^i(k+1)}{\sum_i \alpha_i(k+1)} \end{aligned} \tag{6}$$

Perform linear consensus on the new mean belief

$$\bar{p}^i(k+1) = \bar{p}^i(k) + \varepsilon \sum_{j\in N_i} (\bar{p}^j(k) - \bar{p}^i(k)) \tag{7}$$

Evaluate convergence criterion
end while

Note that this approach directly accounts for the uncertainty in the belief of each agent since the hyperparameters $\alpha^i(k)$ are effectively added at each stage of the iteration. Thus, the solution will be (correctly) biased towards the agent that has the highest counts, and this will be reflected in the calculation of the mean belief.

3.2 Approach # 2: Using a mean-variance approximation

Another way of tackling the problem of probability consensus relies on using the first two moments of the Dirichlet distribution, and approximating the consensus problem using ideas from Kalman Consensus [13]. While this would only be exact for a Gaussian distribution, after a (not very large) number of observations this approximation is fairly reasonable for the Dirichlet. This mean-variance appoach can be applicable if there are existing mean-variance implementations of code, such as a Kalman consensus.

Kalman consensus in the Dirichlet setting relies on the first two moments modulo a minor modification that arises due to the singularity of the Dirichlet covariance matrix. This can be resolved by achieving a consensus on a non-singular submatrix of the Dirichlet covariance matrix, and a component of the mean probability vector in the Mean-variance approximation algorithm (Algorithm 2). The key part of the algorithm is the fact that consensus is only being performed on the submatrix $\Omega^j(k) \doteq (\Sigma^j(k))^{1:M-1\times 1:M-1}$ and on the subvector $\bar{q}^j(k) = (\bar{p}^j(k))^{1:M-1}$. Once the consensus has been reached on the subvector $\bar{q}^i(k)$, the moments must be reconstructed. This can be done very easily because of the unit sum constraint on the mean belief, and the rank deficiency of the full covariance matrix.

Algorithm 2 Mean-variance approximation

Choose a convergence criterion ε
Define: $\bar{q}^i(k) = (\bar{p}^i(k))^{1:M-1}$ and $\Omega^i(k) \doteq (\Sigma^i(k))^{1:M-1\times 1:M-1}$
while Convergence criterion not satisfied **do**
Find next step consensus state:

$$\begin{aligned}
(\Omega^i(k+1))^{-1} &= \sum_{j\in\mathcal{N}_i} \left(\Omega^j(k)\right)^{-1} \\
\bar{q}^i(k+1) &= \bar{q}^i(k) + \Omega^i(k+1) \sum_{j\in N_i} \left(\Omega^j(k)\right)^{-1} (\bar{q}^j(k) - \bar{q}^i(k))
\end{aligned}$$

Evaluate convergence criterion
end while
Upon reaching consensus on $q(k)$, reconstruct the moments (ϕ enforces that the column sum of the covariance matrix is 0):

$$\begin{aligned}
\Sigma^i(k+1) &= \begin{bmatrix} \Omega^i(k+1) & \mathbf{1} - \mathbf{1}^T\Omega^i(k+1) \\ (\mathbf{1} - \mathbf{1}^T\Omega^i(k+1))^T & \phi \end{bmatrix} \\
\bar{p}^i(k+1) &= \left[\bar{q}^i(k+1)^T,\ 1 - \mathbf{1}^T\bar{q}^i(k+1)\right]^T
\end{aligned}$$

The earlier example in Figure 2(b) is now implemented while taking into account the variance information of the Dirichlet and results in a more consistent estimate. There is a good rationalization for why using Kalman consensus is reasonable for the Dirichlet, but for current sake of brevity, is not included. Basically, a good correspondence can be shown to the Beta distribution

4 Results on the agreement problem with information updates

In the previous section, we did not account for active observations by the agents. The problem was primarily one of agreeing over the prior uncertain distributions, but in this section we consider the problem of agreeing over the uncertain probabilities while each agent receives independent observations of the underlying process. In this sense, this problem is closely related to the problem of Distributed Kalman filtering with consensus investigated by Olfati-Saber [15].

4.1 Approach #1A: Dirichlet aggregation with information update

In a modification of the first example, we are now interested in performing consensus in the presence of multinomial-distributed observations. Upon updating their respective Dirichlet distributions, the agents share their updated Dirichlet with their neighbors, in an attempt to agree on the underlying probability. This is done online, so in a sense, the agents are constantly responding to their latest observations, which can be treated as IID with respect to the other agents.

Notation wise we distinguish between temporal updates of the probability in response to actual observations by using the time index t, and maintain the iteration number using the indicator k. The main steps are shown in Algorithm 3 (presented in the next page). Just like in the previous section, linear consensus can still be performed and it properly accounts for the relative confidence of each agent since the calculation of the mean probability (or mean belief) is accounted for using the hyperparameters. At the end of each consensus step, the new Dirichlet at the next time step is the output of the last consensus Dirichlet.

Algorithm 3 Dirichlet Aggregation with information update

Choose a convergence criterion ε
while Convergence criterion not satisfied **do**
Each agent updates the local Dirichlet with local observations (multinomial)

$$f_D(\mathbf{p}_t|\alpha^i(t)) \propto f_D(\mathbf{p}_{t-1}|\alpha^i(t-1)) \times f_M(\mathbf{p}_{t-1}|\beta^i(t) \tag{8}$$

Update own agent probability distribution: $f_D(\mathbf{p}_k|\alpha^i(k)) = f_D(\mathbf{p}_t|\alpha^i(t))$
Update Dirichlet: $f_D(\mathbf{p}_{k+1}|\alpha^i(k+1)) \propto f_D(\mathbf{p}_k|\alpha^i(k)) \times \prod_{j\in N_i} f_D(\mathbf{p}_k|\alpha^j(k))$

Find the mean belief:

$$\begin{aligned} \bar{p}^i(k+1) &= \int_0^1 f_D(\mathbf{p}_{k+1}|\alpha^i(k+1))d\mathbf{p}_{k+1} \\ &= \frac{\alpha^i(k+1)}{\sum_i \alpha_i(k+1)} \end{aligned} \tag{9}$$

Perform linear consensus on the new mean belief

$$\bar{p}^i(k+1) = \bar{p}^i(k) + \varepsilon \sum_{j\in N_i} (\bar{p}^j(k) - \bar{p}^i(k)) \tag{10}$$

Evaluate convergence criterion
end while

New agent distribution is $f_D(\mathbf{p}_t|\alpha^i(t)) = f_D(\mathbf{p}_{k+1}|\alpha^i(k+1))$
Update the time step $t \leftarrow t+1$

4.2 Approach #1B: Mean-variance approximation with information update

We can also extend the mean-variance approach of the previous section by invoking some of our earlier work [18] that demonstrate that it is possible to recursively express the moments of the Dirichlet. This gives rise to the following mean-variance approximation recursive algorithm that is similar in nature to Approach #1A, but uses the first two moments of the Dirichlet instead of the counts $\alpha^i(k)$ (Algorithm [**?**]).

4.3 Performance in simulation, the need for fading

Both Algorithm 3 and Algorithm 4 gives rise to a very rapid aggregation of “counts" (in the consensus stage) or a very rapid decrease in the uncertainty (in the form of the covariance converging rapidly to $\mathbf{0}^{M\times M}$ in the consensus stage) such that in a real-time setting, the observations of the individual agents are not impacted by repeated observations. This “information freeze" is shown in Figure 3(a), where 5 agents attempting to agree on the probability (shown in black), while taking observations of the sample process, simply cannot agree on the probability.

It has been shown in our previous work [18] that the Dirichlet covariance can very quickly converge to $\mathbf{0}^{M\times M}$ with each additional observations, and that this in turn can cause the information update (or innovation term) to be effectively ignored by the moment updates. In a typical Kalman Filter setting, the reason amounts to an inadequate amount of process noise in the estimation process, and this is precisely the issue here: the agents are taking noise free observations, and after only a few small steps, the information from the neighboring agents results in an increased confidence of the probabilities, even though this confidence is unwarranted.

Some of our previous work [18] as shown that one can add an effective process noise to the mean-variance implementation in the form of covariance scaling: namely,

- Process noise addition: $\Sigma^i(k) \leftarrow 1/\lambda\Sigma^i(k)$

Algorithm 4 Mean-variance approximation with information update

Choose a convergence criterion ε

Each agent updates the local Dirichlet with local observations (multinomial)

$$\begin{aligned}\bar{p}_i(t) &= \bar{p}_i(t-1) + \Sigma_{ii}(t-1)\frac{\delta_{i,i'} - \bar{p}_i(t-1)}{\bar{p}_i(t-1)(1-\bar{p}_i(t-1))} \\ \Sigma_{ii}^{-1}(t) &= \gamma_t\, \Sigma_{ii}^{-1}(t-1) + \frac{1}{\bar{p}_i(t-1)(1-\bar{p}_i(t-1))}\end{aligned}$$

Update own agent moments: $\bar{p}^i(k) = \bar{p}^i(t),\ \ \Sigma^i(k) = \Sigma^i(t)$

while Convergence criterion not satisfied **do**

Find next step consensus state:

$$\begin{aligned}(\Omega^i(k+1))^{-1} &= \sum_{j\in\mathcal{N}_i} \left(\Omega^j(k)\right)^{-1} \\ \bar{q}^i(k+1) &= \bar{q}^i(k) + \Omega^i(k+1) \sum_{j\in\mathcal{N}_i} \left(\Omega^j(k)\right)^{-1} (\bar{q}^j(k) - \bar{q}^i(k))\end{aligned}$$

Evaluate convergence criterion

end while

Upon reaching consensus on $q(k)$, reconstruct the moments (ϕ enforces that the column sum of the covariance matrix is 0):

$$\begin{aligned}\Sigma^i(k+1) &= \begin{bmatrix} \Omega^i(k+1) & \mathbf{1} - \mathbf{1}^T\Omega^i(k+1) \\ (\mathbf{1} - \mathbf{1}^T\Omega^i(k+1))^T & \phi \end{bmatrix} \\ \bar{p}^i(k+1) &= \left[\bar{q}^i(k+1)^T,\ 1 - \mathbf{1}^T\bar{q}^i(k+1)\right]^T\end{aligned}$$

New agent moments is $\bar{p}^i(t+1) = \bar{p}^i(k),\ \ \Sigma^i(t+1) = \Sigma^i(k)$

We have good structural justification to show that these amount to performing the identical operation, but the key result is that it can be shown that by performing this fading, we can in fact allow the agents to agree. Ref [18], we showed that this covariance scaling is equivalent to fading the observation counts, and an analogous modification to the counts-based algorithms is to update the hyperparameters as

- Exponentially fading: $\alpha^i(k+1) \leftarrow \lambda\,\alpha^i(k)$

In terms of the covariance matrix, the fading allows the covariance to drive to $\mathbf{0}$ much slower, as shown in Figure 4 for two distinct agents (blue and red), for two different values of $\lambda \in \{0.93, 1\}$ (note the plot is in logarithmic scale). Since the scaled covariance matrix maintains a higher level of uncertainty, this actually allows the agents to ultimately agree to the correct probability (Figure 3(b)).

5 Conclusions

This paper has discussed some notions of agreeing on an uncertain or imprecise probability in a network of agents. We have shown that ignoring the uncertainty in the probability can result in inconsistent estimates, and shown methods by which a consistent estimate can be achieved. More importantly however, we have shown that consensus protocols can result in the agents not converging at all to the desired probability. In essence, the agents must have some form of fading mechanism embedded (either through the application of a covariance scaling or a fading factor) in order to ultimately converge to the desired quantity.

Our ongoing work is investigating the role of distributed consensus on transition probability models for the purposes of robust planning in the context of robust MDPs. Ongoing work is also investigating the impact of heterogeneity in the team (as modeled with each agent having a unique, time-varying scaling or fading factor λ_k) and the question of fault detection in the sensor team, namely the ability to identify if there is a sensor that is broadcasting erroneous information, thereby biasing the team to the wrong probability. Finally, a thorough investigation of what needs to be communicated to the other members of the team (in the spirit of the work of Rosencrantz, Gordon, and Thrun [19]) is being investigated.

6 Acknowledgments

This research was funded in part under Air Force Grants # F49620-01-1-0453 and # FA9550-08-1-0086.

7 References

[1] D. Castañón and D. Teneketzis, "Further results on the consensus problem," in *IEEE Conference on Decision and Control*, 1987.

[2] D. Castañón and D. Teneketzis, "Further results on the asymptotic agreement problem," in *IEEE Transactions on Automatic Control*, vol. 33, pp. 515–523, 1988.

[3] M. DeGroot, "Reaching a consensus," in *Journal of the American Statistical Association*, vol. 69, pp. 118–121, 2008.

[4] S. Fortunato, "The krause-hegselmann consensus model with discrete opinions," in *International Journal of Modern Physics: C*, vol. 15, 2004.

[5] C. Genest and J. Zidek, "Combining probability distributions: A critique and an annotated bibliography," in *Statistical Science*, vol. 1, pp. 114–148, 1986.

[6] A. Josang, "The consensus operator for combining beliefs," in *Artificial Intelligence*, vol. 141, pp. 157–170, 2002.

[7] D. Pennock and M. Welmman, "Graphical representations of consensus belief," in *Proc of Uncertainty in Artificial Intelligence*, 1999.

[8] D. Teneketzis and P. Varaiya, "Consensus in distributed estimation with inconsistent beliefs," in *Systems and Controls Letters*, vol. 4, pp. 217–221, 1984.

[9] R. Winkler, "The consensus of subjective probability distributions," in *Management Science*, vol. 15, pp. 61–75, 1968.

[10] R. Winkler, "Combining probability distributions from dependent information sources," in *Management Science*, vol. 27, pp. 479–488, 1981.

[11] R. Washburn and D. Teneketzis, "Asymptotic agreement among communicating decisionmakers," in *Stochastics*, vol. 13, pp. 103–129, 1984.

[12] J. Pearl, "Probabilistic reasoning in intelligent systems: Networks of plausible inference," in *Morgan-Kauffman*, 1988.

[13] W. Ren, R. W. Beard, and D. B. Kingston, "Multi-agent kalman consensus with relative uncertainty," in *American Control Conference, 2005. Proceedings of the 2005*, pp. 1865–1870 vol. 3, 2005.

[14] W. Ren, R. W. Beard, and E. M. Atkins, "A survey of consensus problems in multi-agent coordination," in *American Control Conference, 2005. Proceedings of the 2005*, pp. 1859–1864 vol. 3, 2005.

[15] R. Olfati-Saber, E. Franco, E. Frazzoli, and J. Shamma, "Belief consensus and distributed hypothesis testing in sensor networks," in *Network Embedded Sensing and Control* (P. Antsaklis and P. Tabuada, eds.), vol. 331 of *Lecture Notes in Computer Science*, pp. 169–182, Springer-Verlag, 2006.

[16] D. P. Spanos, R. Olfati-Saber, and R. M. Murray, "Approximate distributed kalman filtering in sensor networks with quantifiable performance," in *IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks*, (Piscataway, NJ, USA), p. 18, IEEE Press, 2005.

[17] A. Gelman, J. Carlin, H. Stern, and D. Rubin, *Bayesian Data Analysis.* Chapman and Hall, 1995.

[18] L. F. Bertuccelli, "Robust decision making with model uncertainty in aerospace systems," in *PhD Thesis MIT*, 2008.

[19] M. Rosencrantz, G. Gordon, and S. Thrun, "Decentralized sensor fusion with distributed particle filters," in *Proc. 19th Conf. on Uncertainty in Artificial Intelligence (UAI)*, 2003.

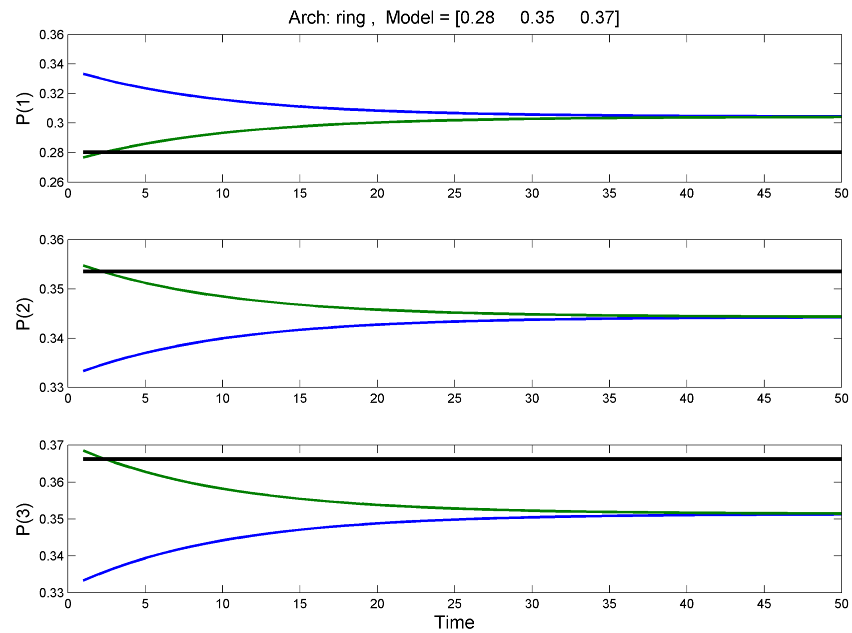

(a) Inconsistent agreement on a probability in a ring network

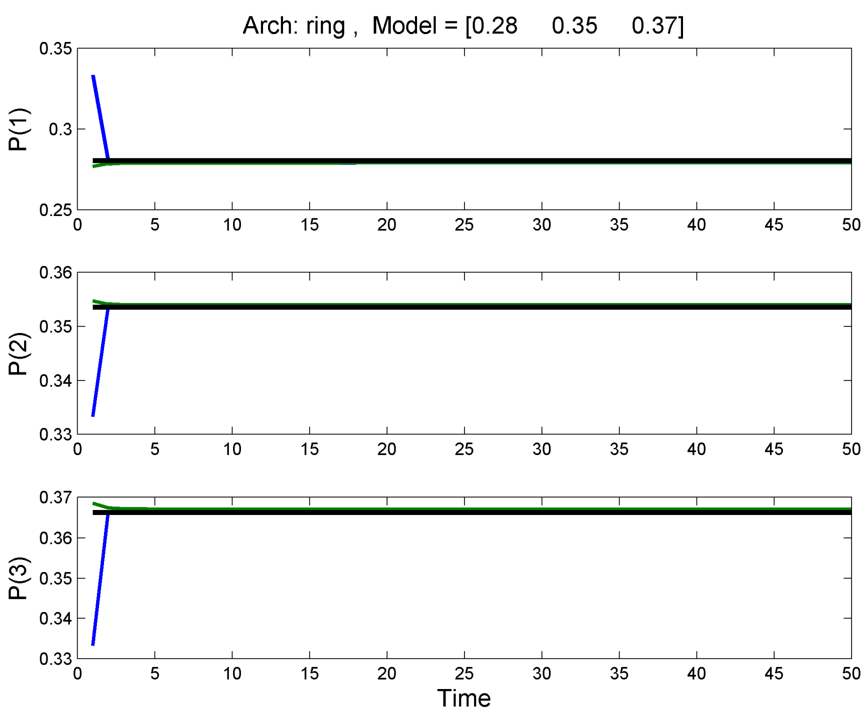

(b) Consistent agreement on a probability in a ring network

Figure 2: Inconsistency in probability estimate

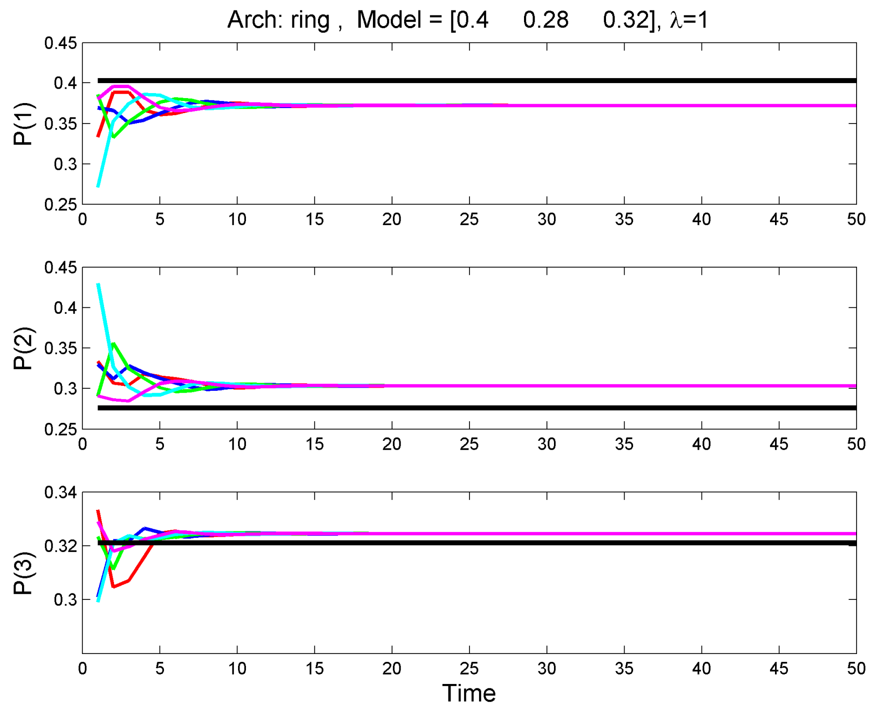

(a) Without any exponential fading (or process noise addition) the agents cannot agree on the probability

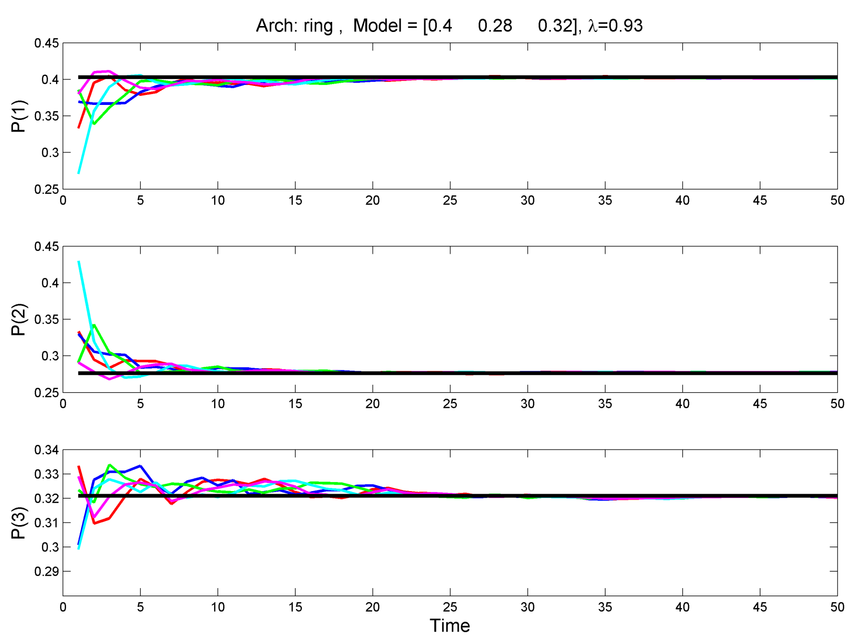

(b) Addition of exponential fading (or process noise addition) allows agents to agree

Figure 3: Importance of fading

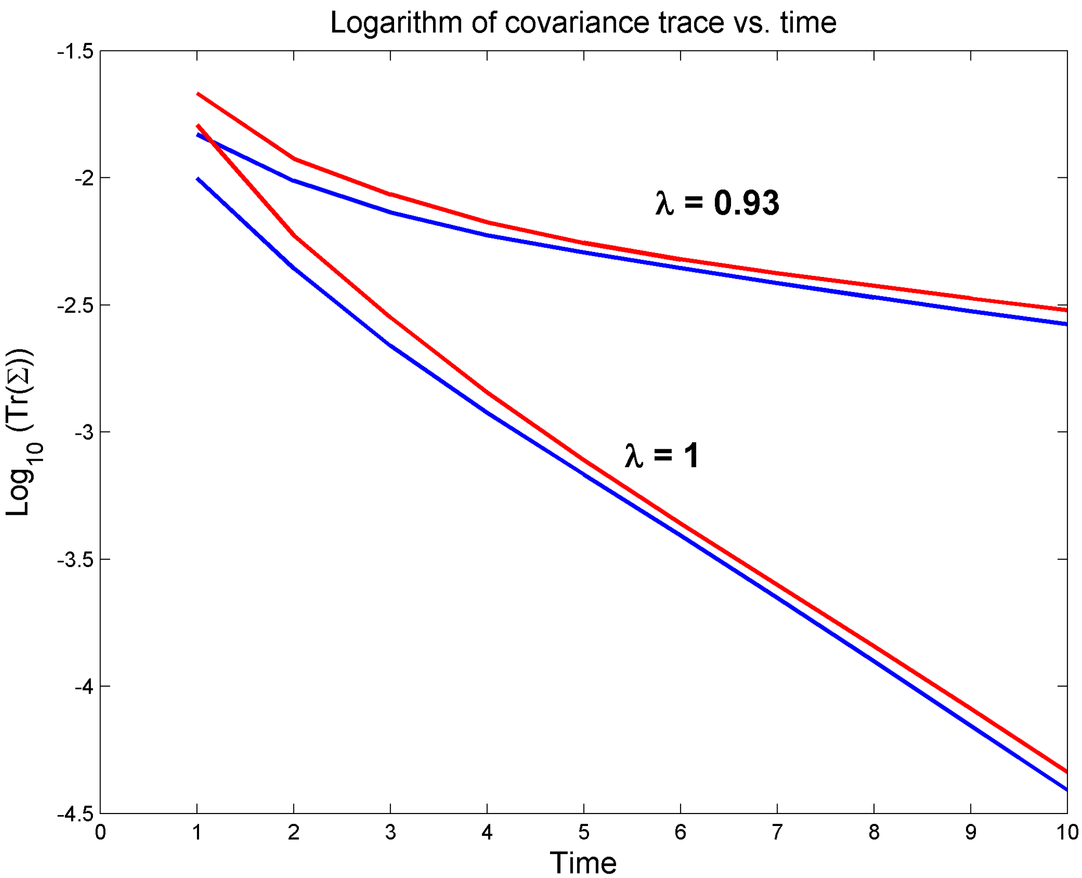

Figure 4: Fading impact on the covariance