CONVERGENCE ANALYSIS OF CONSENSUS BELIEF FUNCTIONS WITHIN ASYNCHRONOUS AD-HOC FUSION NETWORKS

Thanuka L. Wickramrathne[*]

Dept. of Comp. Sci. and Eng. &
Environmental Change Initiative (ECI)
Univ. of Notre Dame, South Bend, IN, USA
t.wickramarathne@nd.edu

Kamal Premaratne, Manohar N. Murthi[†]

Dept. of Elect. and Comp. Eng.
Univ. of Miami, Coral Gables, FL, USA
{kamal,mmurthi}@miami.edu

ABSTRACT

In a multi-agent data fusion scenario, agents may iteratively exchange their states to arrive at a *consensus* state which signifies 'general agreement' among the agents. Agent states that are being exchanged may have been generated from *hard* (i.e., physics based) or *soft* (i.e., human based evidence. such as opinions or beliefs regarding an event) sensors. Convergence analysis becomes an extremely challenging problem in such complex fusion environments, which may involve communication delays, ad-hoc paths, etc. In this paper, we analyze consensus of a Dempster-Shafer theoretic (DST) fusion operator by formulating the consensus problem as finding common fixed points of a pool of paracontracting operators. Due to its DST basis, this consensus protocol can deal with a wider variety of data imperfections characteristic of hard+soft data fusion environments. It also easily adapts itself to networks where agent states are captured with probability mass functions because they can be considered a special case of DST models.

Index Terms— Consensus, data fusion, Dempster-Shafer theory, paracontracting operators, multi-agent systems.

1. INTRODUCTION

Background. Mathematical modeling of consensus appears in [1], and more recently in sensor related research and applications [2–8]. In these works, an agreement is sought among a group of *agents,* where the term "agent" refers to sources, which can be either *soft* (i.e., human-based sensors) or *hard* (i.e., physics-based sensors).

Motivation. In a typical consensus setup, agents iteratively exchange and revise their states until convergence or agreement is reached. In data fusion networks, convergence analysis is a challenging task because of the usually nonlinear iterative processes involved (in combining, updating, and/or revising evidence) and communication related difficulties (e.g., link delays, dynamic and ad-hoc link structure, etc.).

Relation to Prior Work. Consensus analysis within the context of data fusion in distributed sensor networks that takes into account these communication difficulties appears in [2–4, 8]. However, due mainly to the difficulties associated with nonlinear fusion strategies, the majority of these consensus techniques utilize a simple weighted average for belief revision (or state updates) [5, 9].

The work in [10] on nonlinear asynchronous paracontracting operators provides a methodology for convergence analysis involving nonlinear operators under various network topologies with link delays and dynamic link structure. Fang et al. [6] exploits this work to develop asynchronous consensus protocols. This current work of ours also exploits the work in [10] to develop consensus protocols and analyze their convergence properties. However, our work differs from [6] because we use the theoretical underpinnings of [10] to analyze the convergence of certain Dempster-Shafer theory (DST) based data fusion strategies, which have been demonstrated to be more suited for hard+soft fusion scenarios [11–14]. Due to page length limitations, we omit the proofs of our main results in Section 3. These will appear in a future publication. As far as the authors are aware, our work constitutes the first instance where convergence analysis of DST data fusion schemes are being studied. The work in [19] demonstrates how such a consensus may yield valuable information regarding an agent's credibility. By acting as a proxy of a local group of agents, a consensus state may also enable data fusion to be carried out with a significantly lower computational burden. These DST convergence results can be extended to the case in which agents exchange probability mass function (p.m.f.) estimates since p.m.f.s can be considered a special case of DST.

2. PRELIMINARIES

2.1. DS Theory

Basic Notions. In DST, the total set of mutually exclusive and exhaustive propositions of interest (i.e., the 'scope of expertise') is referred to as the *frame of discernment (FoD)* $\Theta = \{\theta_1, \ldots, \theta_n\}$ [15]. A singleton proposition θ_i represents the lowest level of discernible information. Elements in the power set of Θ, 2^Θ, form all the propositions of interest. We use $A \setminus B$ to denote all singletons in A that are not in B; $\overline{A}$ denotes $\Theta \setminus A$.

Definition 1. *Consider the FoD* Θ *and* $A \subseteq \Theta$.

The mapping $m(\boldsymbol{.}) : 2^\Theta \mapsto [0, 1]$ *is a* basic belief assignment (BBA) *or* mass assignment *if* $\sum_{A\subseteq\Theta} m(A) = 1$ *with* $m(\emptyset) = 0$. *The* belief *of* A *is* $Bl(A) = \sum_{B\subseteq A} m(B)$*, and the* plausibility *of* A *is* $Pl(A) = 1 - Bl(\overline{A})$. □

DST captures the notion of *ignorance* by allocating masses to composite propositions (i.e., a non-singleton proposition). A proposition that possesses non-zero mass is a *focal element.* The set of focal elements is the *core* $\mathfrak{F}$*;* the triplet $\mathcal{E} \equiv \{\Theta, \mathfrak{F}, m(\boldsymbol{.})\}$ is the corresponding *body of evidence (BoE).* While $m(A)$ measures the support assigned to proposition A *only,* the belief represents the total

[*]The results reported in this paper are based on work conducted while the first author was at the Department of Electrical and Computer Engineering, University of Miami, Coral Gables, Florida, USA.

[†]This work is based on research supported in part by the U.S. Office of Naval Research (Grants N00014-10-1-0140 and N00014-11-1-0493) and in part by the U.S. National Science Foundation (Grant 1038257).

support that can move into A without any ambiguity; $Pl(A)$ represents the extent to which one finds A plausible. When focal elements are constituted of singletons only, the BBA, belief and plausibility all reduce to a probability assignment.

Data Fusion. Data fusion is the process of combining $\mathcal{E}_i \equiv \{\Theta_i, \mathfrak{F}_i, m_i\}$, $i = 1, 2$, to arrive at a new $\mathcal{E} \equiv \{\Theta, \mathfrak{F}, m\}$ representing the aggregated evidence.

Definition 2 (Dempster's Combination Rule (DCR))**.** *With identical FoDs (i.e., $\Theta_1 = \Theta_2 \equiv \Theta$), the DCR-fused BoE, denoted as $\mathcal{E} \equiv \mathcal{E}_1 \oplus \mathcal{E}_2$, is*

$$m(A) = \sum_{C \cap D = A} m_1(C)\, m_2(D)/(1-K), \, \forall A \subseteq \Theta,$$

whenever $K = \sum_{C \cap D = \emptyset} m_1(C)\, m_2(D) \neq 1$. □

$K \in [0, 1]$ captures the *conflict* between the BoEs being fused. The DCR's difficulties in fusing conflicting BoEs are well documented. The *conditional approach* [11,13,16] offers an elegant way for updating evidence while accommodating potentially contradictory evidence and sources possessing non-identical FoDs.

Definition 3 (Conditional Update Equation (CUE))**.** *[19] For the BoEs $\mathcal{E}_i$, $i = 1, \ldots, n$, the CUE that updates $\mathcal{E}_1$ with the evidence in $\mathcal{E}_i$, $i = 2, \ldots, n$, to generate $\mathcal{E}$, denoted as $\mathcal{E} \equiv \mathcal{E}_1 \lhd (\mathcal{E}_2 \bowtie \cdots \bowtie \mathcal{E}_n)$, is*

$$\mathrm{Bl}(B) = \alpha_1 \ \mathrm{Bl}_1(B) + \sum_{i=2}^{n} \sum_{A \in \mathfrak{F}_i} \beta_{1,i}(A) Bl_i(B|A).$$

Here, the CUE parameters are non-negative and satisfy $\alpha_1 + \sum_{i=2}^{n} \sum_{A \in \mathfrak{F}_2} \beta_{1,2}(A) = 1$. □

The conditional operations in the above definitions are implemented using the Fagin-Halpern (FH) DST conditionals [17], which can be considered the more natural extension of the usual Bayesian conditional notions [14].

2.2. Theory of Paracontractions

Consensus analysis in multi-agent systems can be formulated as a special case of finding common fixed points of a (finite) pool of *paracontracting multiple-point operators* [6]. Convergence of these schemes can then be established if the iterations satisfy certain *coupling* conditions.

Basic Notions. Let $\mathbb{D}$ be the domain of interest (e.g., domain of agent states). A vector $\xi \in \mathbb{D}$ is referred to as a *fixed point* of an operator $F : \mathbb{D}^m \mapsto \mathbb{D}$ iff $F(\xi, \ldots, \xi) = \xi$, where $m \in \mathbb{N}$. Further, the set of all fixed points of operator F is denoted by $fix(F) = \{\xi \in \mathbb{D} \mid F(\xi, \ldots, \xi) = \xi\}$. A vector $\zeta \in \mathbb{D}$ is a *common fixed point* if ζ is a fixed point common to all operators $F \in \mathcal{F}$, viz., $\zeta \in fix(F)$, $\forall F \in \mathcal{F}$.

Let $\mathbb{I} \subset \mathbb{N} \equiv \{1, 2, \ldots\}$ be a set of indices and $m \in \mathbb{N}$ a fixed number. Henceforth, we deal with the *pool of operators* $\mathcal{F} = \{F^i, i \in \mathbb{I} \mid F^i : \mathbb{D}^{m_i} \mapsto \mathbb{D}\}$, where $m_i \in \mathbb{N}$ s.t. $m_i \leq m$, $\forall i \in \mathbb{I}$, and $\mathbb{D}$ is closed.

Definition 4. *[10, 18]*

(i) Paracontracting operator: *An operator $F : \mathbb{D}^m \mapsto \mathbb{D}$ is* paracontracting *on $\mathbb{D}$ w.r.t. a given vector norm $||\,.\,||$, if $\|F(\mathbf{X}) - \xi\| < \max_j \|\mathbf{x}_j - \xi\|$, for all $\mathbf{X} \equiv [\mathbf{x}_1, \ldots, \mathbf{x}_m] \in \mathbb{D}^m$ and any $\xi \in fix(F)$, unless $\mathbf{X} \in fix(F)$.*

(ii) Paracontracting pool of operators: *If for all $i \in \mathbb{I}$, $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_{m_i}] \in \mathbb{D}^{m_i}$ and a vector norm $||\,.\,||$, F^i is continuous on $\mathbb{D}^{m_i}$, then $\mathcal{F}$ is said to be* paracontracting *on $\mathbb{D}$, if for any $\xi \in fix(F^i)$, $\|F^i(\mathbf{X}) - \xi\| < \max_j \|\mathbf{x}_j - \xi\|$, unless $\mathbf{X} \in fix(F^i)$.*

(iii) Asynchronous iteration: *Let $\mathcal{X}_0 = \{\mathbf{x}[-\ell] \in \mathbb{D} \mid \ell = 1, \ldots, M\}$ be a given set of vectors, where M is the number of initial conditions. Let $\mathcal{S}$ denote the sequence of m_i-tuples from $\mathbb{N}_0 \cup \{-1, \ldots, -M\}$, where $\mathcal{S} \equiv \{s^1[k], \ldots, s^{m_{I[k]}}[k] \mid I[k] \in \mathbb{I}, s^\ell[k] \in \mathbb{N}_0 \cup \{-1, \ldots, -M\}$ s.t. $s^\ell[k] \leq k,\ k = 0, 1, \ldots\}$. Then, for sequences $\mathcal{I} \equiv \{I[k] \in \mathbb{I} \mid k = 0, 1, \ldots\}$ and $\mathcal{S}$, the sequence $\mathbf{x}[k+1] = F^{I[k]}\left(\mathbf{x}[\,s^1[k]\,], \ldots, \mathbf{x}[\,s^{m_{I[k]}}[k]\,]\right)$ is referred to as an* asynchronous iteration *and denoted by $(\mathcal{F}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$.* □

Convergence of an asynchronous iteration scheme depends on the properties of the operators and the coupling among the agents.

Definition 5 (Confluent Iteration)**.** *[10] An asynchronous iteration $(\mathcal{F}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$ is* confluent *if there are numbers $n_0, b \in \mathbb{N}$ and a sequence $\{b_k \in \mathbb{N} \mid k = n_0, n_0+1, \ldots$ s.t. $k \geq n_0\}$ s.t. the following is true:*

(i) *for every vertex $k_0 \geq k$, there is a path from b_k to k_0 in $(\mathcal{V}, \mathcal{E})$;*

(ii) *$k - b_k \leq b$;*

(iii) *$\mathcal{S}$ is regulated, i.e., $s \equiv \max_{k,\ell}\left(k - s^\ell[k]\right)$ exists.*

(iv) *for every $i \in \mathcal{I}$, there is a $c_i \in \mathbb{N}$, so that for all $k \geq n_0$, there is a vertex $w_k^i \in \mathcal{V}$, which is a successor of b_k and a predecessor of b_{k+c_i}, for which $I(w_k^i - 1) = i$.* □

The convergence of confluent, nonlinear asynchronous iterations is governed by

Theorem 1. *[10] Let $\mathcal{F}$ be a paracontracting pool on $\mathbb{D} \subset \mathbb{R}^n$ and assume that $\mathcal{F}$ has common fixed points, viz., $fix(\mathcal{F}) \neq \emptyset$. Then, any confluent asynchronous iteration $(\mathcal{F}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$ converges to some fixed point $\xi \in fix(\mathcal{F})$.* □

3. CONSENSUS IN A FUSION ENVIRONMENT

Consider a hard+soft data fusion environment consisting of a set of sensors/agents $\mathcal{N} = \{\mathcal{A}_1, \ldots, \mathcal{A}_m\}$ which interact with each other at discrete time instances $t_0 < t_1 < \cdots < t_k < \cdots$. Here, t_k and k are referred to as the *discrete event-based time* and *discrete event-based time index* [6], respectively. We are now interested in developing a consensus protocol that is applicable to such scenarios. Let us proceed as follows.

3.1. Agent Interactions

Convergence rate and the existence of consensus is dependent on the communication patterns of agents.

Definition 6 (Interaction Topology)**.** Agent interaction topology *refers to the structure of the spatial connectivity among agents at t_k. Let us denote by $Q_{i,j}$, $j = 1, \ldots, n_i$, the j^{th} interaction topology used by agent $\mathcal{A}_i$, for $i = 1, \ldots, m$. Also, let $\mathcal{Q} \equiv \{Q_{i,j} | j = 1, \ldots, n_i;\ i = 1, \ldots, m\}$ be the set of all interaction topologies used by the multi-agent system.* □

An interaction topology is said to be *fully connected* if the agent updating its state receives information from all the other agents. A multi-agent system is fully connected if all the interaction topologies are fully connected. If at least one agent updates its state without taking information from all the other agents, then the corresponding interaction topology is said to be *partially connected*. A multi-agent system is *partially connected* if at least one interaction topology in use is partially connected.

If each agent $\mathcal{A}_i$, $i = 1, \ldots, n$, uses the same interaction topology at time t_k, $k = 0, 1, \ldots$, then such a system is referred to as a *static multi-agent system.* If this does not hold true (i.e., at least one agent uses different interaction topologies), then such a system is referred to as a *dynamic multi-agent system.*

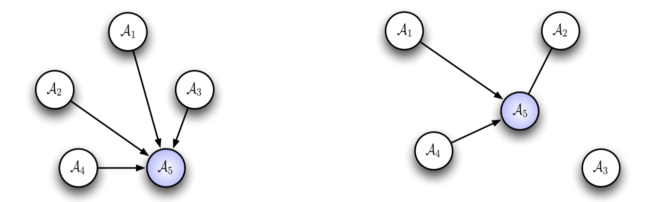

(a) $Q_{5,j}$ is *fully connected*. (b) $Q_{5,j}$ is *partially connected*.

Fig. 1. Spatial connectivity among a set of 5 agents.

3.2. Towards a Rational Consensus

Let us proceed by formally defining consensus.

Definition 7 (Consensus)**.** *Let the state of agent $\mathcal{A}_i$ at k be given by $\mathcal{E}_i[k] \equiv \{\Theta, \mathfrak{F}_i[k], m_i(\centerdot)[k]\}$, for $i = 1, \ldots, m$. Each agent $\mathcal{A}_i \in \mathcal{N}$ starts with an initial state $\mathcal{E}_i[0]$ and repeatedly updates its state (via a valid DST updating strategy) by exchanging information according to some interaction topology. Then, we say that a* consensus *is reached among agents in $\mathcal{N}$, if $\|\mathcal{E}_i[k] - \mathcal{E}_j[k]\| \to 0$ as $k \to \infty$ for all $\mathcal{A}_i, \mathcal{A}_j \in \mathcal{N}$, for some norm $\| \centerdot \|$.* □

The popular *Lehrer-Wagner model* presupposes a weighted average of the agent opinions for generating a consensus [1]. Most consensus protocols and techniques available in the literature espouse this supposition and employ weighted averages of rather simple numerical models (or at most, probability mass functions), and provide only limited utility especially in hard+soft data fusion environments. For instance, in hard+soft fusion environments, one often has access to highly reliable, but vague, estimates of the *ground truth (GT)*. For example, satellite imagery may be used to reliably identify a vehicle as being an SUV, but it may not be able to determine its exact model/year; on the other hand, a witness statement can be more specific but its credibility is questionable.

With these observations in mind, we state several properties that a consensus protocol must satisfy in generating a meaningful consensus, viz., a *rational consensus.* Let us denote the GT via the DST BoE $\mathcal{E}^t_\Theta \equiv \{\Theta, \mathfrak{F}^t_\Theta, m^t_\Theta(\centerdot)\}$, where $\mathcal{F}^t_\Theta$ contains only one singleton proposition (which identifies the GT). Let us a denote a relaible estimate of the GT via the DST BoE $\hat{\mathcal{E}}^t_\Theta \equiv \{\Theta, \hat{\mathfrak{F}}^t_\Theta, \hat{m}^t_\Theta(\centerdot)\}$, where $\hat{\mathfrak{F}}^t_\Theta$ contains only one non-singleton proposition.

Definition 8 (Rational Consensus)**.** *Let $\mathcal{E}^*_\Theta \equiv \{\Theta, \mathfrak{F}^*_\Theta, m^*_\Theta(\centerdot)\}$ denote a consensus reached by the agents in $\mathcal{N}$ via some valid consensus protocol. We say that $\mathcal{E}^*_\Theta$ is* rational *if the following are true.*

(i) *When the GT is $\mathcal{E}^t_\Theta$ is known: $\mathcal{E}^*_\Theta = \mathcal{E}^t_\Theta$; and*

(ii) *When a reliable estimate $\hat{\mathcal{E}}^t_\Theta$ of the GT is known: $\mathcal{E}^*_\Theta$ is a 'refinement' of $\hat{\mathcal{E}}^t_\Theta$, i.e., $\forall B \in \mathfrak{F}^*_\Theta$, $\exists C \in \hat{\mathfrak{F}}^t_\Theta$ s.t. $B \subseteq C$.* □

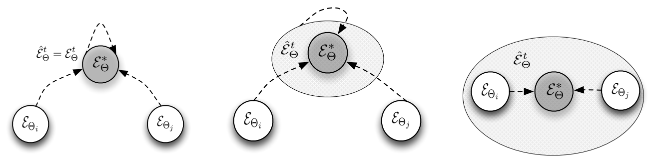

(a) GT known. (b) GT estimate known. (c) Unknown GT.

Fig. 2. Pictorial illustration of a DST rational consensus.

We now show that the DST evidence updating strategy CUE in [11] generates a rational consensus in the sense of Definition 8. This further strengthens the suitability of the CUE in hard+soft fusion environments because of its functional definition (which can be given a probabilistic interpretation), flexibility in parameter selection, and most importantly, robustness against contradictory evidence and agents possessing non-identical expertise.

3.3. Consensus Protocols for DST BoEs

Consider the CUE with its parameters selected as in

Definition 9. *Let $\mathfrak{E}_\Theta \equiv \{\mathcal{E} \mid \mathcal{E} = \{\Theta, \mathfrak{F}, m(\centerdot)\}\}$ denote the set of all possible BoEs defined on Θ. Now, consider the set of n BoEs, $\mathcal{E}_i \in \mathfrak{E}_\Theta,\ i = 1, \ldots, n$. Then, the operator $\mathrm{F}_\triangleleft{}^{ij} : \mathfrak{E}^n_\Theta \mapsto \mathfrak{E}_\Theta$ that updates $\mathcal{E}_i$ with respect to interaction topology $Q_{i,j}$ (updating w.r.t $\mathcal{E}_{j_1}, \ldots, \mathcal{E}_{j_n}$) is defined as*

$$\mathrm{F}_\triangleleft{}^{ij}(\mathcal{E}_{j_1}, \ldots, \mathcal{E}_{j_n}) \equiv \mathcal{E}_i \triangleleft (\mathcal{E}_{j_1} \bowtie \cdots \bowtie \mathcal{E}_{j_n}),$$

where the CUE parameters are given by

$$\alpha_i = C_i, \quad \beta_{ij_k}(A) = \begin{cases} C_{j_k}\, m_i(A), & \text{for } \mathcal{E}_i \equiv \hat{\mathcal{E}}^t; \\ C_{j_k}\, m_{j_k}(A), & \text{otherwise}, \end{cases}$$

s.t. $\alpha_i + \sum_{j_1,\ldots,j_n} \sum_{A \in \mathfrak{F}_{j_k}} \beta_{ij_k}(A) = 1$, *where C_i is positive real.* □

Claim 2. *The operator $\mathrm{F}_\triangleleft{}^{ij} : \mathfrak{E}^n_\Theta \mapsto \mathfrak{E}_\Theta$ in Definition 9 is paracontractive on $\mathfrak{E}_\Theta$ w.r.t. any p-norm.* ■

This leads us to the following DST consensus protocol.

Definition 10 (Asynchronous DST Consensus Protocol)**.** *An asynchronous consensus protocol generated by the operator $\mathrm{F}_\triangleleft{}^{ij}$ is*

$$\mathcal{E}_i[k+1] = \begin{cases} \mathrm{F}_\triangleleft{}^{ij}\left(\mathcal{E}_{j_1}[s^{j_1}(k)], \ldots, \mathcal{E}_{j_n}[s^{j_n}(k)]\right), & I[k] = Q_{ij}; \\ \mathcal{E}_i[k], & \textit{otherwise}. \end{cases}$$

This corresponds to the asynchronous iteration $(\mathcal{F}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$, where $\mathcal{I} = \{I[k] \in \mathcal{Q}\}$, $\mathcal{F} = \{\mathrm{F}_\triangleleft{}^{ij} \mid Q_{ij} \in \mathcal{I}\}$, $\mathcal{X}_0 = \{\mathcal{E}_i[0],\ i = 1, \ldots, n\}$, and $\mathcal{S}$ s.t. $s^\ell[k] \in \mathbb{N}_0$ s.t. $s^\ell[k] \leq k$ for $\ell = 1, \ldots, m,\ k = 0, 1, \ldots$ □

This protocol allows for CUE-based evidence updating of agent BoEs in an asynchronous fashion, where agent interaction topologies can be either fully or partially connected, and either static or dynamic. Indeed, a very important results that justifies the use of this protocol in hard+soft fusion environments is

Claim 3. *A consensus BoE generated using the protocol in Definition 10 is a rational consensus in accordance with Definition 8.* □

We now make use of Theorem 1 to establish convergence. Due to page length limitations, the proofs of these results are omitted.

3.4. Convergence Analysis

Synchronous Fully Connected Topology. This represents perhaps the simplest agent setup where each agent is connected to all the other agents and information is exchanged without any iteration delay (i.e., $k - s^j[k] = 0$). In this case, consensus protocol in Definition 10 reduces to

$$\mathcal{E}_i[k+1] = \mathrm{F}_\triangleleft{}^{i}(\mathcal{E}_1[k], \ldots, \mathcal{E}_m[k]),\ i \in \overline{1,m},\ k \geq 1. \quad (1)$$

This iteration is confluent and hence the protocol converges.

Synchronous, Static, Partially Connected Topology. This represents an agent setup where agents communicate without iteration delays with a network topology that is not fully connected and that does not change over time. Therefore, not all agents can communicate to all the others. In this case, the consensus protocol in Definition 10 reduces to

$$\mathcal{E}_i[k+1] = \mathrm{F}_{\triangleleft}{}^{i} \left(\mathcal{E}_1[k], \ldots, \mathcal{E}_{m_i}[k]\right), \tag{2}$$

for $i = 1, \ldots, m$, and $k \geq 1$, where $m_i = 1, \ldots, m$. This iteration converges as long as the *graph union* of interaction networks of all agents is connected.

Synchronous, Dynamic, Partially Connected Network. This represents an agent setup where agents communicate without iteration delays with a network topology that is not fully connected and changes over time. Then, the consensus protocol in Definition 10 reduces to

$$\mathcal{E}_{\Theta_i}[k+1] = \mathrm{F}_{\triangleleft}{}^{I[k]} \left(\mathcal{E}_{\Theta_1}[k], \ldots, \mathcal{E}_{\Theta_{m_{I[k]}}}[k]\right), \tag{3}$$

for $i = 1, \ldots, m$, and $k \geq 1$, where $m_i = 1, \ldots, m$.

Given that $\mathcal{F}_{\triangleleft}$ is paracontracting on $\mathfrak{E}_{\Theta}$, we only need to satisfy confluence conditions on the equivalent iteration $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \mathcal{S})$ to prove convergence. In a synchronous network, assumption (iii) of Definition 5, viz., $k - s^{\ell}(k) \leq s, \forall k \in \mathbb{N}_0,\ \ell = 1, \ldots, m$, for any $s \in \mathbb{N}_0$ is clearly satisfied. Assumption (i) of Definition 5 is also clearly satisfied. Thus, if the interaction sequence $\mathcal{I}$ is regulated (i.e., one should be able to find a finite time span at any given time t_k, on which all the agents participate in the consensus process), while the interaction graphs $G[k]$, $k \geq 1$, of the network are repeatedly jointly rooted, the consensus protocol converges.

Asynchronous Fully Connected Network. This represents an agent setup where each agent is connected to all the other agents, but the information exchange is not synchronized (or delayed) (i.e., $k - s^j[k] < 0$). In this case, the consensus protocol in Definition 10 reduces to

$$\mathcal{E}_{\Theta_i}[k+1] = \mathrm{F}_{\triangleleft}{}^{i} \left(\mathcal{E}_{\Theta_1}[s^1(k)], \ldots, \mathcal{E}_{\Theta_m}[s^m(k)]\right), \tag{4}$$

for $i = 1, \ldots, m$, and $k \geq 1$.

The convergence of this consensus protocol occurs if the iteration delays are finite.

This discussion can also be extended to study asynchronous, static and partially connected networks. In this case, similar to the case of static, partially connected networks, one needs to impose conditions on how agents interact in order to ensure that there is adequate coupling among agents to reach a consensus. This can be achieved by making sure that graph union of interaction topologies of each agent is a connected graph. In order to guarantee the confluence in iterations, one also needs to ensure that the iteration delays are finite at any given time.

4. NUMERICAL EXAMPLE

Consider a fusion scenario [19] that consists of 5 interacting agents whose initial beliefs/opinions are given by the following BoEs $\mathcal{E}_i[0]$, $i = 1, \ldots, 5$, with $\Theta_i \equiv \Theta = \{abcde\}$:

BBA	a	c	d	e	ac	ad	abc	bc	cd	cde
$m_1(\centerdot)[0]$	0.8	-	0.1	-	-	-	-	0.1	-	-
$m_2(\centerdot)[0]$	-	0.9	-	-	-	-	-	-	0.1	-
$m_3(\centerdot)[0]$	-	-	0.1	-	-	-	0.9	-	-	-
$m_4(\centerdot)[0]$	-	-	-	0.1	0.9	-	-	-	-	
$m_5(\centerdot)[0]$	-	-	-	-	-	0.9	-	-	-	0.1

Let us now look at the convergence of this multi-agent system under the following configurations.

Figure 3 shows the variation of distance $\|\mathcal{E}_1[k] - \mathcal{E}_2[k]\|$ for each configuration as the system evolves.

Remarks:

(a) The parameters α and C_i were chosen to be equal for all the experiments.

Config.	**Static/ Dynamic**	**Fully/Partially Connected**	**Delays**
1	static	fully	no (synchronous)
2	static	partially (graph union connected)	no (synchronous)
3	static	partially (graph union not connected)	no (synchronous)
4	static	partially	finite
5	dynamic	partially	finite

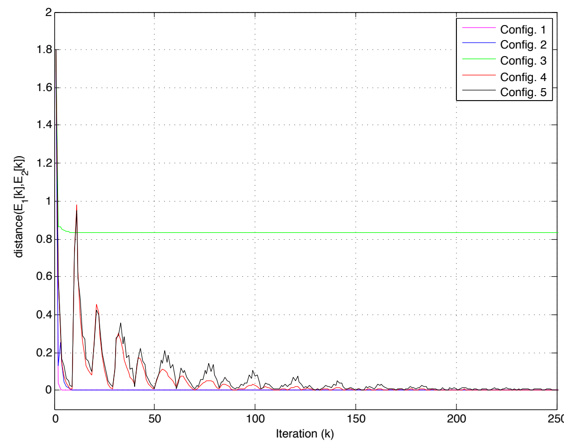

Fig. 3. Consensus formation (as seen from $\|\mathcal{E}_1[k] - \mathcal{E}_2[k]\|$).

(b) Config. 3 illustrates that, if the graph union of partially connected topologies is not connected, then no formation of consensus occurs.

(c) We looked at the finite delay case only, because networks with infinite delays do not converge, in general.

(d) Networks without delays (Configs. 1 and 2) converge much faster with less oscillatory behavior.

5. CONCLUSION

This work develops a DST protocol that allows evidence sources in a multi-agent environment to iteratively exchange their states to arrive at a consensus state. This consensus protocol is well suited for use in hard+soft data environments where the types and variety of imperfections call for evidence models that are based on DST. Capitalizing on results regarding fixed points of a pool of paracontracting multiple-point operators, the proposed protocol is shown to generate a consensus under a wide variety of network topologies (in the presence or absence of communication delays) and under static and dynamic conditions.

6. REFERENCES

[1] K. Lehrer and C. Wagner, *Rational Consensus in Science and Society*, ser. Philosophical Studies Series in Philosophy. D. Reidel Pub. Co., 1981.

[2] R. Olfati-Saber and J. S. Shamma, "Consensus filters for sensor networks and distributed sensor fusion," in *IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, Dec. 2005, pp. 6698 – 6703.

[3] L. Xiao, S. Boyd, and S. Lall, “A scheme for robust distributed sensor fusion based on average consensus,” in *International Symposium on Information Processing in Sensor Networks (IPSN)*, Apr. 2005, pp. 63 – 70.

[4] L. Schenato and G. Gamba, “A distributed consensus protocol for clock synchronization in wireless sensor network,” in *IEEE Conference on Decision and Control*, Dec. 2007, pp. 2289 – 2294.

[5] M. Cao, A. S. Morse, and B. D. O. Anderson, “Reaching a consensus in a dynamically changing environment: A graphical approach,” *SIAM Journal of Control and Optimization*, vol. 47, no. 2, pp. 575–600, 2008.

[6] L. Fang and P. J. Antsaklis, “Asynchronous consensus protocols using nonlinear paracontractions theory,” *IEEE Transactions on Automatic Control*, vol. 53, no. 10, pp. 2351 –2355, Nov 2008.

[7] S. Kar and J. M. F. Moura, “Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise,” *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 355 –369, jan 2009.

[8] Z. Wu, H. Fang, and Y. She, “Weighted average prediction for improving consensus performance of second-order delayed multi-agent systems,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, to appear, 2012.

[9] F. Borran, M. Hutle, N. Santos, and A. Schiper, “Quantitative analysis of consensus algorithms,” *IEEE Transactions on Dependable and Secure Computing*, vol. 9, no. 2, pp. 236 –249, Mar. 2012.

[10] M. Pott, “On the convergence of asynchronous iteration methods for nonlinear paracontractions and consistent linear systems,” *Linear Algebra and its Applications*, vol. 283, no. 13, pp. 1 – 33, 1998.

[11] K. Premaratne, M. N. Murthi, J. Zhang, M. Scheutz, and P. H. Bauer, “A Dempster-Shafer theoretic conditional approach to evidence updating for fusion of hard and soft data,” in *Proc. International Conference on Information Fusion (ICIF)*, Seattle, WA, July 2009, pp. 2122–2129.

[12] T. L. Wickramarathne, K. Premaratne, M. N. Murthi, and M. Scheutz, “A Dempster-Shafer theoretic evidence updating strategy for non-identical frames of discernment,” in *Proc. Workshop on the Theory of Belief Functions (WTBF’10)*, Brest, France, Apr. 2010.

[13] T. L. Wickramarathne, K. Premaratne, M. N. Murthi, M. Scheutz, and S. Kübler, “Belief theoretic methods for soft and hard data fusion,” in *Proc. International Conference on Statistical Signal Processing (ICASSP)*, Prague, Czech Republic, May 2011.

[14] T. L. Wickramarathne, K. Premaratne, and M. N. Murthi, “Toward efficient computation of the dempster-shafer belief theoretic conditionals,” *IEEE Transactions on Systems, Man and Cybernetics—Part B: Cybernetics*, to appear, 2012.

[15] G. Shafer, *A Mathematical Theory of Evidence*. Princeton, NJ: Princeton University Press, 1976.

[16] E. C. Kulasekere, K. Premaratne, D. A. Dewasurendra, M.-L. Shyu, and P. H. Bauer, “Conditioning and updating evidence,” *International Journal of Approximate Reasoning*, vol. 36, no. 1, pp. 75–108, Apr. 2004.

[17] R. Fagin and J. Y. Halpern, “A new approach to updating beliefs,” in *Proc. Conference on Uncertainty in Artificial Intelligence (UAI)*, P. P. Bonissone, M. Henrion, L. N. Kanal, and J. F. Lemmer, Eds. New York, NY: Elsevier Science, 1991, pp. 347–374.

[18] M. Pott, “On the convergence of asynchronous iteration methods for nonlinear paracontractions and consistent linear systems,” in *Preprint: Sonderforschungsbereich Diskrete Strukturen in der Mathematik*, vol. 95, no. 32, 1995.

[19] T. Wickramarathne, K. Premaratne, and M. N. Murthi, “Consensus-based credibility estimation of soft evidence for robust data fusion,” in *Belief Functions*, ser. Advances in Intelligent and Soft Computing, T. Denoeux and M.-H. Masson, Eds. Compiégne, France: Springer Berlin / Heidelberg, May 2012, vol. 164, pp. 301–309.