Convergence Analysis of Iterated Belief Revision in Complex Fusion Environments

Thanuka L. Wickramarathne, *Member, IEEE*, Kamal Premaratne, *Senior Member, IEEE*, Manohar N. Murthi, *Member, IEEE*, and Nitesh V. Chawla, *Member, IEEE*

***Abstract*—We study convergence of iterated belief revision in *complex fusion environments*, which may consist of a network of *soft* (i.e., human or human-based) and *hard* (i.e., conventional physics-based) sensors and where agent communications may be asynchronous and the link structure may be dynamic. In particular, we study the problem in which network agents exchange and revise *belief functions* (which generalize *probability mass functions*) and are more geared towards handling the uncertainty pervasive in soft/hard fusion environments. We focus on belief revision in which agents utilize a generalized fusion rule that is capable of generating a *rational consensus*. It includes the widely used *weighted average consensus* as a special case. By establishing this fusion scheme as a pool of *paracontracting operators*, we derive general convergence criteria that are relevant for a wide range of applications. Furthermore, we analyze the conditions for consensus for various social networks by simulating several network topologies and communication patterns that are characteristic of such networks.**

***Index Terms*—Belief revision, conditional update equation, consensus, convergence of opinions, soft/hard fusion.**

I. Introduction

UNDERSTANDING the collective behavior of a group of adaptive *agents* (i.e., sensors, fusion nodes, people, etc.), where such behavior is solely governed by the *local interactions*, is a topic of overlapping interest for many researchers in engineering, mathematics, and the physical and social sciences [1]–[19]. To quote [3], which succinctly captures the essence of such collective behavior of a flock of geese flying in tight formation,

> . . . collectively they [geese] form the image of a giant delta-shaped bird that moves as purposively as if it were a single organism. Yet the flock has no 'group mind' nor is there a 'leader bird' choreographing the formation. Rather each bird reacts to its immediate neighbors who in turn react to it.

In reaching a *consensus*, which can also be viewed as a special case of collective group behavior, a group of "rational" agents iteratively interact (or communicate) with their neighbors and update their beliefs (or precisely, the *agent state*) toward collectively estimating some phenomenon of interest without any global coordination. Key questions of concern in consensus analysis include, the type(s) of interactions that will eventually lead to a consensus, convergence of a given class of interactions, rate of convergence, etc.

Motivation: Sensing methods are increasingly moving toward distributed architectures where large numbers of heterogeneous sources are used for information acquisition. These sources can include a mix of both *soft* (i.e., human or human-based) and *hard* (i.e., conventional physics-based) sensors that are embedded in (often) *unstructured* environments.[1] For instance, *participatory sensing* exploits the wide-array of sensors in consumer mobile devices for gathering local knowledge [20]. Consensus analysis then can be viewed as the study of the convergence behavior of an iterated belief revision process that is embedded in a *complex fusion* network.[2] With the emergence of social networks as a dominant force in modern society, understanding the mathematical underpinnings that may lead to a consensus (e.g., of opinions) in such complex fusion environments is crucial from the perspectives of signal processing, economics, sociology, political science, and many other fields of study.

Challenges: Convergence analysis in complex fusion networks is a challenging task mainly due to the unstructured environment, and the difficulties associated with modeling the (often uncertain) agent state. Most consensus studies are undertaken in binary or multi-alternative decision spaces, where the agent state is modeled via real-valued vectors [4]–[18]. Such modeling cannot adequately well capture the types of uncertainties and the nuances that are characteristic of agent states in complex fusion environments. Imprecise probabilistic

Manuscript received September 30, 2013; revised January 27, 2014; accepted March 14, 2014. Date of publication April 21, 2014; date of current version July 16, 2014. This work was supported by the University of Notre Dame Faculty Research Support Program for T. L. Wickramarathne and the U.S. Office of Naval Research via Grants #N00014-10-1-0140 and #N00014-11-1-0493 and the U.S. National Science Foundation via Grant #1038257 for K. Premaratne, M. N. Murthi, and T. L. Wickramarathne (while he was at the University of Miami), and the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053 for NVC. The guest editor coordinating the review of this manuscript and approving it for publication was Prof. Vikram Krishnamurthy.

T. L. Wickramarathne and N. V. Chawla are with the Interdisciplinary Center for Network Science and Applications, Environmental Change Initiative, and the Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46556 USA (e-mail: t.wickramarathne@nd.edu; nchawla@nd.edu).

K. Premaratne and M. N. Murthi are with the Department of Electrical and Computer Engineering, University of Miami, Coral Gables, FL 33146 USA (e-mail: kamal@miami.edu; mmurthi@miami.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSTSP.2014.2314854

[1]We use the term *unstructured network/environment* to refer to ad-hoc/dynamic network topologies with possibly *asynchronous* communication (e.g., delays) between agents.

[2]We use the term *complex fusion network/environment* to refer to a collection of soft/hard sensors which employ possibly nonlinear fusion operations that are embedded in an unstructured network.

formalisms, such as *Dempster-Shafer (DS) theory* [21], are better suited for handling these data uncertainties (see Table II), but convergence properties of such formalisms in unstructured environments have yet to be explored.

Contributions: In this paper, we explore the convergence of iterated belief revision of a certain DS theoretic (DST) fusion operator [22]–[26] functioning within a complex fusion network. As far as the authors are aware, our work constitutes the first instance where convergence of DST data fusion schemes is studied in complex fusion networks. By exploiting the convergence notions of *asynchronous iterations of paracontracting operators* [27], we devise belief revision strategies that can generate what we refer to as a *rational consensus* (see Claim 8). This rational consensus guarantees a more meaningful consensus state that, for example, is consistent with a reliable estimate of the ground truth. We are not aware of other consensus protocols (e.g., the weighted average) having the ability to provide such a rational consensus. In this sense too, this current work is novel.

Relevance: The DS framework is better suited for modeling multiple types of data uncertainties that are characteristic of soft/hard fusion environments (see Table II). It also bears a close relationship to the probabilistic framework: (**a**) the inner and outer measures of a non-measurable event turn out to be closely related to the DS theoretic (DST) notions of *belief* and *plausibility*, respectively [28], [29]; and (**b**) when positive support is only assigned to *singleton propositions* (see Section II-B), the belief and plausibility functions become identical to yield *probability mass functions (p.m.f.s)*. Therefore, belief/plausibility functions can be viewed as generalizations of p.m.f.s, or conversely, p.m.f.s can be considered as a limiting case of belief/plausibility functions. Hence, by deriving convergence criteria for a generalized DST fusion operator to generate a rational consensus in an asynchronous and dynamic sensor network, we essentially provide a mechanism to explore both DST and probabilistic data fusion strategies in one formalism (e.g., probabilistic weighted average consensus where agents exchange p.m.f. estimates). Therefore, especially for researchers who are primarily working in probabilistic frameworks, the attraction of this DST approach is immediately apparent: (**a**) it allows for a smoother transition to probabilistic notions which after all are quite effective and hence well entrenched in work related to sensor fusion; and (**b**) while engendering the use of belief functions that generalize p.m.f.s, it provides a more intuitive means to capture uncertainty by allowing support to be allocated to propositions that do not consist solely of singletons.

Our work has tremendous relevance to problems in distributed data fusion and opinion dynamics in social networks. Given the greater ability of belief functions to handle uncertainty and ambiguous evidence, they provide a better model for an agent's opinion about a topic or situation, or a fusion node's assessment of a situation. 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. Furthermore, as the work in [30] demonstrates, a rational consensus can be utilized to estimate agent credibility, when the ground truth is known only partially or even unknown.

Relation to Prior Work: Much of the previous work in consensus has focused on applications in distributed control, estimation, or fusion [4]–[8], [10]–[12], [14], [16]–[18]. In control problems, the agents communicate among each other over a network to achieve a consensus for a control objective. In estimation problems, the agents collectively attempt to estimate an underlying vector or statistic (e.g., mean) of a signal, often within the framework of statistical signal processing (e.g., Kalman filtering). In fusion, the agents collectively pool their evidence to arrive at a consensus decision. In these works, the agents update their states (usually in a distributed manner) using a linear (e.g., weighted average) or nonlinear interaction model; the interactions network may have communication impairments; some papers also take into account noisy observations (i.e., incoming evidence). But most (if not all) of these previous works deal with an agent state that is modeled as a real-valued vector, i.e., the consensus state is a real-valued vector. In our work, the agent state is a DST body of evidence (BoE) (which subsumes a p.m.f.). In contrast to other models that are perhaps better suited for control and estimation purposes, a DST agent state model is ideally suited for capturing the data imperfections and nuances associated with soft data (e.g., witness statements, domain expert opinions, etc.). However, the interaction protocols that we can use are severely restricted because *each entry in the agent state must be non-negative and all entries must sum to one*, a constraint that is absent in previous works. With this restricted set of interaction protocols, achieving consensus is necessarily more challenging in our work: the protocol not only has to generate a converging agent state but it also must ensure that the agent state being generated at *each* iteration remains a DST BoE (or a p.m.f. as the case may be).

Moreover, the rational consensus notion that our protocol achieves is different than what appears in existing literature, where a rational consensus typically means a weighted average [1]. As we claim later, this latter notion is not adequate for working in complex fusion environments. In contrast, our notion of a rational consensus is more meaningful in that the consensus we achieve has a very clear relationship with the ground truth, an aspect that is absent in the previous works.

The work in [27] on *nonlinear asynchronous paracontracting operators* provides the mathematical framework for convergence analysis of *asynchronous iterations* involving both finite and infinite pools of nonlinear operators. Fang, *et al.* [31] exploit this work to develop asynchronous *consensus protocols* for finite pools of *convex combination operators* [32]. While we take inspiration from the work in [31] to view consensus protocols as asynchronous iterations, our work differs from [31] due to several reasons. We use the theoretical underpinnings of [27] to analyze the convergence of certain DST data fusion operators defined via *conditional belief functions* of agent states. Due to the fact that DST belief functions are defined as Möbius transformations on DST *mass assignments* (see Section II-B), identifying whether such fusion schemes map their inputs to the relative interior of the convex hull of inputs (which would allow for direct application of [31]) is extremely difficult. Furthermore, we also use convergence results on infinite pools in [27] to derive conditions under which a consensus protocol consisting of infinite number of fusion operators converge. Our previous work in [26] provide only a summary of the convergence results (without proofs) of a finite pool of DST fusion operators.

TABLE I
SUMMARY OF NOTATION

Notation	Description
Θ, 2^Θ	Frame of discernment (FoD), power set of Θ.
$\lvert B \rvert$	Cardinality of set B.
$A \setminus B$	Set difference of A and B, i.e., $\{x \in A \text{ s.t. } x \notin B\}$.
$\overline{B}$	Set complement of $B \subseteq \Theta$, i.e., $\overline{B} \equiv \Theta \setminus B$.
$m(\bullet)$, $\mathrm{Bl}(\bullet)$, $\mathrm{Pl}(\bullet)$	DST notions of mass, belief, and plausibility. The corresponding conditional notions are $m(\bullet \mid A)$, $\mathrm{Bl}(\bullet \mid A)$, $\mathrm{Pl}(\bullet \mid A)$, where $A \subseteq \Theta$ s.t. $\mathrm{Bl}(A) > 0$ is the conditioning set.
$\mathbb{D}$	Domain of all probabilistic assignments on Θ.
$\mathbb{N}$, $\mathbb{N}_0$	Set of positive integers, set of non-negative integers.
$\Re$, $\Re_+$, $\Re_{+0}$	Set of reals, set of positive reals, set of non-negative reals.
$\mathbb{I}$	Finite index set from the positive integers, i.e., $\mathbb{I} \subset \mathbb{N}$.
$\mathcal{A}_i$	Agent $i \in \{1, \ldots, m\}$, where $m \in \mathbb{N}$ is fixed.
$\mathcal{F}$	Pool of operators $F^i : \mathbb{D}^{m_i} \mapsto \mathbb{D}$ s.t. $m_i \leq m$, $\forall i \in \mathbb{I}$.
$i \leftarrow I[k]$	Agent $\mathcal{A}_i$ identified via $F^{I[k]}$.
t_k	discrete event-based time, where $k \in \mathbb{N}_0$.
$\mathcal{I} = \{I[k]\}_{k=0}^{\infty}$	Sequence identifying the operator $F^{I[k]} \in \mathcal{F}$ that is active at t_k.
$\mathcal{S} = \{S[k]\}_{k=0}^{\infty}$	Sequence of m_i-tuples $S[k] = (s^1[k], \ldots, s^{m_{I[k]}}[k])$.
GT, $\widehat{GT}$	Ground truth, estimate of ground truth

This paper is organized as follows: Section II provides a review of essential notions; Section III describes DST modeling of belief revision; Section IV provides a perspective where consensus is viewed as a fixed-point of multiple-point operators; Section V presents the convergence analysis of belief revision, and provides conditions under which agents embedded in several popular network topologies converge to a consensus; Section VI provides several experimental studies; finally, Section VII provides concluding remarks. The proofs are relegated to the Appendix for clarity of the presentation. Table I summarizes the notation being employed.

II. Preliminaries

A. Graph Theory Notions for Modeling Agent Interactions

Definition 1 (Basic Notions): [10], [33] A **graph** $G = (V, E)$ consists of a **vertex set** $V \subset \mathbb{N} \equiv \{1, 2, \ldots\}$, an **edge set** $E \subseteq V \times V$, and a relation that associates each edge with two vertices. A **directed graph** or **digraph** is a graph where E is a set of ordered pairs of vertices. An ordered pair $(u, v) \in E$, or $u \to v$, denotes a directional edge from $u \in V$ to $v \in V$. For $u \to v$, u and v are the **predecessor** of v and **successor** of u, respectively. A vertex $v \in V$ is a **root** of G, if $\exists u \to v$, $\forall u \in V, u \neq v$. A **rooted graph** has at least one root. □

Definition 2 (Graph Composition): [10] Consider arbitrary graphs $G[k] = (V[k], E[k]) \in \mathbb{G}$, where $V[k] \subset \mathbb{N}$ for $k = 1$, 2, where $\mathbb{G}$ is the set of all graphs with vertex set $\subset \mathbb{N}$. Then, the **composition** $G[2] \circ G[1]$ is the graph s.t. the edge set contains edges (u, w), if $(u, v) \in E[1]$ and $(v, w) \in E[2]$.

A finite sequence of directed graphs $(G[k] \in \mathbb{G})_{k=1}^{\ell}, \ell \in \mathbb{N}$ is **jointly rooted**, if $G[\ell] \circ G[\ell - 1] \circ \cdots \circ G[1]$ is a rooted graph. An infinite sequence of graphs $(G[k] \in \mathbb{G})_{k=1}^{\infty}$ is a **repeatedly jointly rooted**, if $\exists \ell \in \mathbb{N}$ for which each finite sequence $(G[i])_{i=\ell(k-1)+1}^{k\ell}$, $k \geq 1$ is jointly rooted. □

B. DST Framework for Modeling Agent States

1) Basic Notions: In DS theory, the total set of mutually exclusive and exhaustive propositions of interest (i.e., *sample space*) is referred to as the *frame of discernment (FoD)* $\Theta = \{\theta_1, \ldots, \theta_n\}$ [21]. A *singleton proposition* $\theta_i \in \Theta$ represents the lowest level of discernible information. Elements in 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 3: Consider the FoD Θ and $B \subseteq \Theta$. The mapping $m(\cdot) : 2^\Theta \mapsto [0, 1]$ is a **basic probability assignment (BPA)** or **mass assignment** if $\sum_{B \subseteq \Theta} m(B) = 1$ with $m(\emptyset) = 0$. The **belief** and **plausibility** of B are $Bl(B) = \sum_{C \subseteq B} m(C)$ and $Pl(B) = 1 - Bl(\overline{B})$, respectively. A proposition $B \subseteq \Theta$ that possesses non-zero mass is a **focal element**. The set of focal elements is the **core** $\mathcal{F}$; the triplet $\mathcal{E} \equiv \{\Theta, \mathcal{F}, m(\cdot)\}$ is the corresponding **body of evidence (BoE)**. The set of all BoEs defined on Θ is denoted by $\mathcal{E}_\Theta = \{\mathcal{E} \mid \mathcal{E} = \{\Theta, \mathcal{F}, m(\cdot)\}\}$. □

DST supports modeling of *ignorance* by allowing partial probability specifications via direct assignment of non-zero mass to *composite* (in contrast to singletons) propositions. While $m(B)$ measures the support that is directly assigned to $B \subseteq \Theta$ *only*, the belief $\mathrm{Bl}(B)$ represents the total support that can move into B without any ambiguity; $\mathrm{Pl}(B)$ represents the extent to which one finds B plausible. These DST belief and plausibility measures are closely related to the inner and outer measures of a non-measurable event $B \subseteq \Theta$ w.r.t. p.m.f.s defined on Θ. Furthermore, when focal elements are constituted of singletons only, the mass, belief and plausibility all reduce to a p.m.f..

2) Conditional Notions: The Fagin-Halpern conditionals extend the notions of inner and outer measures of non-measurable events to conditioning.

Theorem 1 (Fagin-Halpern (FH) Conditionals): [29] For any $B \subseteq \Theta$ and a conditioning event $A \subseteq \Theta$ s.t. $\mathrm{Bl}(A) > 0$, conditional belief $\mathrm{Bl}(B|A) : 2^\Theta \mapsto [0, 1]$ and conditional plausibility $\mathrm{Pl}(B|A) : 2^\Theta \mapsto [0, 1]$ are given by $\mathrm{Bl}(B|A) = \mathrm{Bl}(A \cap B)/[\mathrm{Bl}(A \cap B) + \mathrm{Pl}(A \cap \overline{B})]$ *and* $\mathrm{Pl}(B|A) = \mathrm{Pl}(A \cap B)/[\mathrm{Pl}(A \cap B) + \mathrm{Bl}(A \cap \overline{B})]$, respectively.

A proposition with positive mass after conditioning is referred to as a *conditional focal element*. The collection of conditional focal elements that are generated with respect to the conditioning event A is referred to as the *conditional core* and denoted by $\mathcal{F}_{|A}$. Thus, $\mathcal{F}_{|A} = \{B \subseteq \Theta | m(B|A) > 0\}$ and $m(\cdot|A) : 2^\Theta \mapsto [0, 1]$ is the corresponding conditional BPA related to $\mathrm{Bl}(\cdot|A)$ via the Möbius transformation [21]

$$m(B|A) = \sum_{C \subseteq B} (-1)^{|B-C|} \mathrm{Bl}(C|A), \quad \forall B \subseteq \Theta. \tag{1}$$

An important recent result that can be used to directly identify the FH conditional focal elements is

Theorem 2 (Conditional Core Theorem (CCT)): [25] Given $\mathrm{Bl}(A) > 0$ in the BoE $\mathcal{E} = \{\Theta, \mathcal{F}, m(\cdot)\}$, $m(B|A) > 0$ iff B can be expressed as $B = X \cup Y$, for some $X \in \mathrm{in}(A)$ and $Y \in \mathrm{OUT}(A) \cup \{\emptyset\}$. Here, $\mathrm{in}(A) = \{B \subseteq A | B \in \mathcal{F}\}$ and $\mathrm{OUT}(A) = \{B \subseteq A | B = \bigcup_{i \in \mathcal{I}} C_i, C_i \in \mathrm{out}(A)\}$, where $\mathrm{out}(A) = \{B \subseteq A | B \cup C \in \mathcal{F}, \emptyset \neq B, \emptyset \neq C \subseteq \overline{A}\}$. ■

The work presented in this paper employs FH conditional notions for belief revision. For a comprehensive discussion on the CCT, we refer the interested reader to [25]. The following example [34] illustrates the application of the CCT.

Example 1: [34] Consider the BoE, $\mathcal{E} = \{\Theta, \mathcal{F}, m(\cdot)\}$ with $\Theta = \{a, b, c, d, e, f, g, h, i\}$, $\mathcal{F} = \{a, b, h, df, beg, \Theta\}$ and $m(B) = \{0.1, 0.1, 0.1, 0.2, 0.2, 0.3\}$, for $B \in \mathcal{F}$ (in the same order given in $\mathcal{F}$). Then, for $A = (abcde)$, $\text{in}(A) = \{a, b\}$; $\text{out}(A) = \{d, be, abcde\}$; and $\text{OUT}(A) = \{d, be, bde, abcde\}$. So, according to the CCT, the conditional focal elements (w.r.t. A) are $\{a, b, ad, bd, be, abe, bde, abde, abcde\}$, because these are the only propositions that can be expressed as $X \cup Y$, where $X \in \text{in}(A)$ and $Y \in \text{out}(A) \cup \{\emptyset\}$. ■

C. Paracontracting Operators for Establishing Convergence

Let $\mathbb{D} \subset \Re^{|2^\Theta|}$ be the domain of interest (e.g., all probabilistic assignments on the sample space Θ). A vector $\xi \in \mathbb{D}$ is referred to as a **fixed point** of an operator $F : \mathbb{D}^m \mapsto \mathbb{D}$, if $F(\xi, \ldots, \xi) = \xi$, where $m \in \mathbb{N} \equiv \{1, 2, \ldots\}$. Further, the set of all fixed points of operator F is denoted by $fix(F) \equiv \{\xi \in \mathbb{D} | F(\xi, \ldots, \xi) = \xi\}$. A vector $\zeta \in \mathbb{D}$ is a *common fixed point* if ζ is a fixed point of each operator $F \in \mathcal{F}$, i.e., $\zeta \in fix(F)$, $\forall F \in \mathcal{F}$. Let $\mathbb{I} \subset \mathbb{N}$ be a set of indices and $m \in \mathbb{N}$ a fixed number. Henceforth, we deal with the **pool of operators** $\mathcal{F} \equiv \{F^i : \mathbb{D}^{m_i} \mapsto \mathbb{D} | m \geq m_i \in \mathbb{N}, \forall i \in \mathbb{I}\}$, where $\mathbb{D}$ is closed. Let $\mathbf{X} \equiv [\mathbf{x}_1, \ldots, \mathbf{x}_{m_i}] \in \mathbb{D}^{m_i}$ and $\mathbf{Y} \equiv [\mathbf{y}_1, \ldots, \mathbf{y}_{m_i}] \in \mathbb{D}^{m_i}$ be arbitrary vectors on $\mathbb{D}^{m_i}$. Also, let $\mathbb{N}_0 \equiv \{0\} \cup \mathbb{N}$ and $\|\cdot\|$ be a vector norm defined on $\mathbb{D}$.

Definition 4 (Paracontracting Operators). [27]:

(i) The pool $\mathcal{F}$ is **contractive** on $\mathbb{D}$, if $\|F^i(\mathbf{X}) - F^i(\mathbf{Y})\| \leq \omega \cdot \max_j \|\mathbf{x}_j - \mathbf{y}_j\|$, $\forall i \in \mathbb{I}, \mathbf{X}, \mathbf{Y} \in \mathbb{D}^{m_i}$, for some $0 \leq \omega < 1$.

(ii) An operator $F^i \in \mathcal{F}$ is **paracontracting** on $\mathbb{D}$ if $\|F^i(\mathbf{X}) - \xi\| < \max_j \|\mathbf{x}_j - \xi\|$, $\forall \mathbf{X} \in \mathbb{D}^{m_i}$ and any $\xi \in fix(F^i)$, unless $\mathbf{X} \in fix(F^i)$.

(iii) Let F^i be continuous on $\mathbb{D}^{m_i}$ for all $i \in \mathbb{I}$. Then, the pool $\mathcal{F}$ is said to be **paracontracting** on $\mathbb{D}$, if $\|F^i(\mathbf{X}) - \xi\| < \max_j \|\mathbf{x}_j - \xi\|$ for any $\xi \in fix(F^i)$, unless $\mathbf{X} \in fix(F^i)$.

□

Definition 5: [27] An infinite pool of operators $\mathcal{G} \equiv \{G^k : \mathbb{D}^{\tilde{m}_k} \mapsto \mathbb{D} | \tilde{m}_k \in \{1, \ldots, m\}, \forall k \in \mathbb{N}\}$ is said to **approximate** the pool $\mathcal{F}$, if there exists $i_k \in \mathbb{I}$, $\forall k \in \mathbb{N}_0$, s.t. $\tilde{m}_k = m_{i_k}$, for which $\lim_{j\to\infty} \|G^j(\cdot) - F^{i_j}(\cdot)\| = 0$, $\forall j \in \mathbb{N}_0$, uniformly for all $\mathbf{X} \in \mathbb{D}^m$. □

We view consensus in terms of notions of paracontracting operators by establishing fusion operators as paracontractions, and taking their fixed points as consensus states. Asynchronous iterations are used for modeling iterative belief revision.

Definition 6 (Asynchronous Iteration): [27] Let $\mathcal{X}_0 = \{\mathbf{x}[-\ell] \in \mathbb{D} | \ell = 1, \ldots, M\}$ be a set of initial conditions, where $M \in \mathbb{N}_0$. Let $\mathcal{S}$ denote the sequence of m_i-tuples of the form $(s^1[k], \ldots, s^{m_{I[k]}}[k])$, where $s^\ell[k] \in \mathbb{N}_0 \cup \{-1, \ldots, -M\}$ s.t. $s^\ell[k] \leq k$ for all $\ell \in \{1, \ldots, m_{I[k]}\}$, where $I[k] \in \mathbb{I}$, $\forall k \in \mathbb{N}_0$. Then, for sequences $\mathcal{I} \equiv \{I[k] \in \mathbb{I} | k \in \mathbb{N}_0\}$ and $\mathcal{S}$, the sequence $\mathbf{x}[k+1] = F^{I[k]}(\mathbf{x}[s^1[k]], \ldots, \mathbf{x}[s^{m_{I[k]}}[k]])$ is referred to as an **asynchronous iteration** and is denoted by $(\mathcal{F}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$. □

Definition 7: [27] Let $(\mathcal{F}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$ be an asynchronous iteration. Then,

(i) $\mathcal{S}$ is **admissible**, if $s^\ell[k] \to \infty$, $\forall \ell$, for $k \to \infty$;

(ii) $\mathcal{S}$ is **regulated**, if $s \equiv \max_{k,\ell}(k - s^\ell[k])$ exists;

(iii) $\mathcal{I}$ is **admissible**, if $I[k] \cup I[k+1] \cup \cdots = \mathbb{I}$, $\forall k \in \mathbb{N}_0$;

(iv) $\mathcal{I}$ is an **index-regulated sequence** if, for all $i \in \mathbb{I}$, $\exists c_i \in \mathbb{N}_0$, s.t. $i \in \{\{I[k]\} \cup \{I[k+1]\} \cup \cdots \cup \{I[k+c_i]\}\}$, $\forall k \in \mathbb{N}_0$;

(v) $\mathcal{I}$ is **regulated** if $\exists c \in \mathbb{N}_0$ s.t. $I[k] \cup I[k+1] \cup \cdots \cup I[k+c] = \mathbb{I}$, $k \in \mathbb{N}_0$.

□

Remark: Clearly, $\mathbb{I}$ has to be finite for $\mathcal{I}$ to be regulated.

Operator coupling for convergence analysis can be studied by associating a directed graph with asynchronous iterations [27]: every iteration including initial conditions are assigned a *vertex*; a pair (j, k) is an *edge* of the graph iff the j-th iteration is used for the k-th iteration.

Definition 8 (Confluent Iteration): [27] Let $(\mathcal{F}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$ be an asynchronous iteration. Then the **graph** of $(\mathcal{F}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$ is the directed graph $G \equiv (V, E)$, where the vertex set V is given by $\mathbb{N}_0 \cup \{-1, \ldots, -M\}$ and the edge set E is given by $(k, k_0) \in E$ iff $\exists 1 \leq \ell \leq m_{I[k_0-1]}$ s.t. $s^\ell[k_0 - 1] = k$.

Then, $(\mathcal{F}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$ is called **confluent** if there are numbers $n_0, b \in \mathbb{N}$ and a sequence $\{b_k \in \mathbb{N} | k = n_0, n_0 + 1, \ldots \text{s.t.} k \geq n_0\}$ s.t. the following are true: **(i)** for every vertex $k_0 \geq k$, $\exists b_k \to k_0$ in G; **(ii)** $k - b_k \leq b$; **(iii)** $\mathcal{S}$ is regulated; and **(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 $I(w_k^i - 1) = i$. □

A confluent iteration establishes the necessary conditions for adequate coupling among agents in order to be able to reach a consensus (see Section III-D). The main convergence result in [27] on the convergence of confluent asynchronous iterations is stated in [31] as follows:

Theorem 3: [27], [31] Let $\mathcal{F}$ be a finite paracontracting pool on $\mathbb{D}$ and assume that $\mathcal{F}$ has common fixed points. Then, any confluent asynchronous iteration $(\mathcal{F}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$ converges to some fixed point $\xi \in fix(\mathcal{F})$. □

Theorem 4: [27] Let $\mathcal{F}$ be a finite contractive pool on $\mathbb{D}$ w.r.t. some norm $\|\cdot\|$. Assume $\mathcal{F}$ has a common fixed point $\xi \in \mathbb{D}$, and that the pool $\mathcal{G}$ (which is possibly infinite) approximates $\mathcal{F}$. Then, the asynchronous iteration $(\mathcal{G}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$ is well-defined and converges to $\xi \in \text{fix}(\mathcal{F})$ for initial conditions $\mathcal{X}_0 \subset \mathcal{D}_0$, where

$$\mathcal{D}_0 = \left\{ x \in \mathbb{D} | \quad \|y - \xi\| \leq \|x - \xi\| + \frac{\bar{\epsilon}_{max}}{1 - \omega}, \quad \forall y \in \mathbb{D} \right\}.$$

Here, $\bar{\epsilon}_{max} = \max_{j\geq 1} \sum_{\ell=p_t}^{p_{t+1}-1} \epsilon_j$, $\forall j \in \mathbb{N}$, with $p_t = \min\{p \in \mathbb{N}_0 | s^\ell[j] > p_{t-1}, \forall j \geq p_t, \ell = 1, \ldots, m_{k_j}, \forall t \in \mathbb{N}$, where $p_0 = -M - 1$. □

III. Consensus in a Complex Fusion Environment

The collective behavior of a multi-agent system is usually characterized by individual agents having access to mostly *imperfect* (i.e., incomplete, uncertain, etc.) and decentralized information, absence of global control, communication impairments, etc [3]. Let us now present a DST modeling approach for a unified analysis of convergence phenomena as applicable to a wide array of complex fusion environments.

A. Problem Statement

Consider the FoD $\Theta \equiv \{\theta_1, \ldots, \theta_n\}$. Let $\mathbb{D}$ denote the domain of all probabilistic (including partial) assignments on Θ. Now, consider a multi-agent system $\mathcal{N} \equiv \{\mathcal{A}_1, \ldots, \mathcal{A}_m\}$

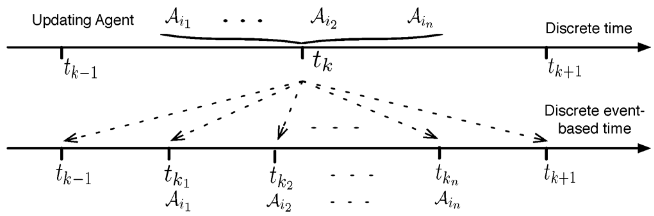

Fig. 1. **Discrete event-based time** refers to event-based time, not necessarily to an absolute time reference. In a situation where n agents $\{\mathcal{A}_{i_1}, \ldots, \mathcal{A}_{i_n}\} \subseteq \mathcal{N}$ update their states at 'absolute' time t_k, a discrete event-based time sequence $\cdots < t_{k-1} < t_{k_1} < t_{k_2} < \cdots \leq t_{k_n} < t_{k+1} \cdots$ can still be extracted s.t. only agent $\mathcal{A}_{i_j}$ updates its state at time t_{k_j} s.t. $k_j \in \mathbb{N}_0, j = 1, \ldots, n$.

consisting of a fixed number of $m \in \mathbb{N}$ adaptive agents interacting toward collectively estimating some phenomenon $\mathcal{X}$, whose "truth" lies in the discrete set Θ of alternatives. The estimate maintained by $\mathcal{A}_i \in \mathcal{N}$ at time t_k is referred to as the *state of* $\mathcal{A}_i$ *at* t_k and we denote it as $\mathbf{x}_i[k] \in \mathbb{D}$; $t_k, k \in \mathbb{N}_0$ denote the *discrete event-based time* [31].

Let $\mathcal{X}_0 \equiv \{\mathbf{x}_i[0] \in \mathbb{D} | i = 1, \ldots, m\}$ denote the set of initial conditions of $\mathcal{N}$, where $\mathbf{x}_i[0]$ denote the initial estimate of $\mathcal{X}$ generated by $\mathcal{A}_i \in \mathcal{N}$. The agents in $\mathcal{N}$ iteratively exchange information with their *neighbors* ($\subset \mathcal{N}$) and update their states via some fusion operator $F \in \mathcal{F}_{\triangleleft} \equiv \{F^i : \mathbb{D}^{m_i} \mapsto \mathbb{D} | m \geq m_i \in \mathbb{N}, i \in \mathbb{I}\}$. Here, $\mathbb{I} \subset \mathbb{N}$ is an index set and $\mathcal{F}_{\triangleleft}$ denotes the pool of fusion operators used by all agents throughout the entire fusion process. Furthermore, an operator $F \in \mathcal{F}_{\triangleleft}$ uniquely identifies the agent $\mathcal{A}_i \in \mathcal{N}$ and its neighbors that are being used for belief revision via F. We are interested in two problems:

[**P1**] Determine a class of fusion operators $\mathcal{F}_{\triangleleft}$ that can (**a**) carry out belief revision robustly, and (**b**) generate a *rational* (i.e., meaningful) agreement among agents; and

[**P2**] Characterize agent interactions that will eventually lead to a consensus among agents.

B. State Updates in a Multi-Agent System

Note that the only time $\mathcal{N}$ undergoes changes is when some agent $\mathcal{A}_i \in \mathcal{N}$ updates its state. Accordingly, $t_k, k \in \mathbb{N}_0$ refer to an event-based time, not necessarily to an absolute time reference. Therefore, without loss of generality, we assume that only one agent updates its state at any given t_k, and the sequence $\{t_k\}_{k=0}^{\infty}$ satisfies $t_k < t_{k+1}, k \in \mathbb{N}_0$ (see Fig. 1).

With the above in place, let us now formalize the iterative belief revision process that takes place in a multi-agent system.

Definition 9: Let $\mathcal{I} = \{I[k] \in \mathbb{I}\}_{k=0}^{\infty}$ denote the sequence of indices identifying the fusion operator $F^{I[k]} \in \mathcal{F}_{\triangleleft}$ that is "active" at $t_k, k \in \mathbb{N}_0$. Let $i \leftarrow I[k]$ denote the mapping $\mathcal{F} \mapsto \mathcal{N}$ identifying the agent $\mathcal{A}_i \in \mathcal{N}$ that is being updated via $F^{I[k]}$ at $t_k, k \in \mathbb{N}_0$. Then, the belief revision process that takes place at t_k can be expressed as:

$$\mathbf{x}_j[k+1] = \begin{cases} F^{I[k]}\left(\mathbf{x}_{i_1}\left[s^1[k]\right], \right. \\ \qquad \left. \ldots, \mathbf{x}_{i_{m_{I[k]}}}\left[s^{m_{I[k]}}[k]\right]\right), & j = i \leftarrow I[k]; \\ \mathbf{x}_j[k], & j \neq i. \end{cases}$$

Here, $\mathcal{S} = \{S[k]\}_{k=0}^{\infty}$ denotes a sequence of m_i-tuples $S[k] \equiv (s^1[k], \ldots, s^{m_{I[k]}}[k])$ with $s^{\ell}[k] \in \mathbb{N}_0 \cup \{-1, \ldots, -m\}$ s.t. $s^{\ell}[k] \leq k$, $\forall \ell \in \{1, \ldots, m_{I[k]}\}, m_i \leq m$, $\forall i \in \mathbb{I}$. □

Remark: Note that at a given time t_k, only one agent state gets updated and states of all other agents remain the same. $I[k]$ identifies the operator $F^{I[k]} \in \mathcal{F}$ performing the state update at t_k, which in turn uniquely identifies the agent $\mathcal{A}_i, i \leftarrow I[k]$ that is getting updated with the information received from its set of neighbors $\mathcal{N}_{I[k]} \subseteq \mathcal{N}$.

Example 2: Consider the multi-agent system $\mathcal{N} = \{\mathcal{A}_1, \ldots, \mathcal{A}_5\}$. Suppose the agents update their states via the pool of operators $\mathcal{F} \equiv \{F^i | i = 1, \ldots, 6\}$, where $F^1 : \mathbb{D}^2 \mapsto \mathbb{D}$, $F^2 : \mathbb{D}^2 \mapsto \mathbb{D}$, $F^3 : \mathbb{D}^4 \mapsto \mathbb{D}$, $F^4 : \mathbb{D}^3 \mapsto \mathbb{D}$, $F^5 : \mathbb{D}^2 \mapsto \mathbb{D}$ and $F^6 : \mathbb{D}^3 \mapsto \mathbb{D}$. Take the initial conditions of $\mathcal{N}$ as $\mathcal{X}_0 = \{\mathbf{x}_i[0] | i = 1, \ldots, 5\}$. Consider the following sequence of state updates:

$$\begin{aligned} \mathbf{x}_1[1] &= F^1(\mathbf{x}_1[0], \mathbf{x}_3[0]), \quad k = 0; \\ \mathbf{x}_2[2] &= F^2(\mathbf{x}_2[1], \mathbf{x}_3[0]), \quad k = 1; \\ \mathbf{x}_4[3] &= F^4(\mathbf{x}_3[0], \mathbf{x}_4[0], \mathbf{x}_5[0]), \quad k = 2; \\ \mathbf{x}_5[4] &= F^5(\mathbf{x}_4[3], \mathbf{x}_5[3]), \quad k = 3; \\ \mathbf{x}_3[5] &= F^3(\mathbf{x}_1[1], \mathbf{x}_2[2], \mathbf{x}_3[4], \mathbf{x}_4[3]), \quad k = 4; \\ \mathbf{x}_5[6] &= F^6(\mathbf{x}_1[1], \mathbf{x}_4[3], \mathbf{x}_5[4]), \quad k = 5; \\ \mathbf{x}_1[7] &= F^1(\mathbf{x}_1[1], \mathbf{x}_3[5]), \quad k = 6, \ etc. \end{aligned}$$

Note that we have omitted the cases for which $\mathbf{x}_j[k]$ is not updated, for which $\mathbf{x}_j[k+1] = \mathbf{x}_j[k], k \geq 0$. By defining the sequences $\mathcal{S} \equiv \{(s^1[0] = 0, s^2[0] = 0), (s^1[1] = 0, s^2[1] = 0), (s^1[2] = 0, s^2[2] = 0, s^3[2] = 0), \ldots\}$ of m_i-tuples, and $\mathcal{I} \equiv \{1, 2, 4, 5, 3, 6, 1, \ldots\}$ of elements of $\mathbb{I} \equiv \{1, \ldots, 6\}$, the belief revision process can be expressed as given in Definition 9. ■

Definition 9 describes a very general belief revision process. It can be modeled via $\mathcal{F}_{\triangleleft}$ and the two sequences $\mathcal{I}$ and $\mathcal{S}$. In fact, our analysis is based on devising a pool of fusion operators that is applicable for soft/hard fusion environments and characterizing $\mathcal{I}$ and $\mathcal{S}$ under which a belief revision process of the type in Definition 9 converges.

C. Uncertainty Modeling for Agent States and Belief Revision

To understand convergence of belief revision processes in complex fusion environments, we utilize DST notions for modeling agent states and the belief revision process.

1) State Models: Let us model $\mathbf{x}_i[k]$ via a DST BoE $\mathcal{E}_i[k] \equiv \{\Theta, \mathcal{F}_i[k], m_i(\cdot)[k]\}, k \in \mathbb{N}_0$, for $i = 1, \ldots, m$. Thus, $\mathbf{x}_i[k]$ can be expressed via any valid mass assignment on 2^{Θ}. This allows for a convenient representation of a wide variety of data imperfections, including probabilistic uncertainties. See Table II. Further, since the DST model is completely specified by mass, belief, or plausibility values, depending on the application, $\mathbf{x}_i[k]$ can be equivalently represented via $m_i(\cdot)[k], \mathrm{Bl}_i(\cdot)[k]$, or $\mathrm{Pl}_i(\cdot)[k]$.

2) Belief Revision: To be able to study a wider array of application domains, we model belief revision via the *conditional update equation (CUE)* proposed in [22], which provides a versatile belief revision mechanism that is applicable to multiple fusion scenarios.

Definition 10 (Conditional Update Equation (CUE)): The CUE that updates $\mathcal{E}_i[k] \in \mathcal{E}_{\Theta}$ with the evidence in $\mathcal{E}_j[k] \in \mathcal{E}_{\Theta}, j \in \{1, \ldots, m\} \setminus \{i\}$, is denoted as $\mathcal{E}_i[k] \triangleleft_{j=1}^{m} \mathcal{E}_j[k]$, and the corresponding belief function is given by

$$\mathrm{Bl}_i(B)[k+1] = \alpha_i[k]\mathrm{Bl}_i(B)[k] + \sum_{j \neq i} \sum_{A \in \mathcal{F}_j[k]} \beta_{ij}(A)[k]\mathrm{Bl}_j(B|A)[k],$$

TABLE II
DST MODELS FOR VARIOUS TYPES OF DATA IMPERFECTIONS [35]

Type of Imperfection	$\mathbf{x}_i[k]$ Focal Set	$m_i(\bullet)[k]$	Remarks
Hard evidence	θ_1	1.0	one singleton foal element
Probabilistic	θ_1	0.1	singleton focal
	θ_2	0.7	elements only
	θ_3	0.2	
Possibilistic	θ_1	0.7	*consonant* focal
	(θ_1, θ_3)	0.2	elements only
	Θ	0.1	
Ambiguity	(θ_1, θ_2)	1.0	inability to discern
Vacuous	Θ	1.0	missing/unknown entry
DST	$\sum_{B\subseteq\Theta} m_i(B)[k] = 1.0$		encapsulates all above and more general forms

where the CUE parameters $\alpha_i[k], \beta_j(A)[k] \in \Re_{+0}$ satisfy

$$\alpha_i[k] + \sum_{j\neq i} \sum_{A\in\mathcal{F}_j[k]} \beta_{ij}(A)[k] = 1, \quad \forall k \in \mathbb{N}_0.$$

Here, $\sum_{j\neq i}$ denotes the sum over $j \in \{1, \ldots, m\} \setminus \{i\}$. □

CUE, which is a convex sum formulation of conditionals, provides a robust mechanism for handling contradictory evidence, vacuous evidence, evidence sources possessing non-identical scopes of expertise, etc; uncertainty of the fused result given by $Un(\cdot)[k] = \mathrm{Pl}_i(\cdot)[k] - \mathrm{Bl}_i(\cdot)[k]$ provides valuable information regarding the underlying uncertainties of the fused result [22], [24], [36]. Moreover, the convex sum formulation is also consistent with the most widely used weighted average view of consensus analysis [1].

3) Characteristics of Agent Interactions: The CUE parameters $\alpha_i[k]$ and $\beta_{ij}(A)[k]$ provide the flexibility to model different belief revision scenarios. As we illustrate in the sections to follow, convergence of multiple fusion scenarios can be analyzed in a unified setting via an appropriate parameterization. We now provide some insight into CUE parameter selection from a consensus perspective and refer the interested reader to [22], [36] for a detailed discussion.

a) Selection of $\alpha_i[k]$: Parameter $\alpha_i[k] \in [0, 1]$ can capture the flexibility of $\mathcal{E}_i[k]$ towards changes, perhaps depending on the credibility of the incoming evidence and/or "inertia" of the available evidence. For instance, a lower $\alpha_i[k]$ is suitable in the initial phase of evidence collection, when $\mathcal{E}_i[k]$ has little or no credible knowledge base to begin with. Another strategy is *inertia-based selection* [36], where both existing and incoming evidence are considered to be equally important.

b) Selection of $\beta_{ij}(\cdot)[k]$: Parameters $\beta_{ij}(\cdot)[k]$ are used for both "refining" and "weighing" incoming evidence. Here, we immediately see two choices [22].

(b.1) Receptive CUE (rCUE): Choose $\beta_{ij}(A)[k] = C_{ij}[k]m_j(A)[k], C_{ij}[k] \in \Re_+$ for all $A \in \mathcal{F}_j[k]$ s.t. the condition in Definition 10 is satisfied. The rCUE strategy weighs incoming evidence based on the support that already exists in the source providing the evidence for each focal element, which can be interpreted as the source receiving evidence being *receptive* to incoming evidence. Not surprisingly, rCUE has an interesting Bayesian interpretation: it yields a weighted average of the corresponding p.m.f.s (see Section III-C-4).

(b.2) Cautious CUE (cCUE): Choose $\beta_{ij}(A)[k] = C_{ij}[k]m_i(A)[k], C_{ij}[k] \in \Re_+$ for all $A \in \mathcal{F}_j[k]$ s.t. the condition in Definition 10 is satisfied. The cCUE strategy weighs incoming evidence based on the support that already exists in the source receiving the evidence for each focal element, which can be interpreted as the source receiving evidence being *cautious* to incoming evidence. Therefore, the focal elements of the updated BoE are restricted to within itself and only refinements are allowed, i.e., a new focal element $B \in \mathcal{F}_i[k+1]$ is created only if there exists a $C \in \mathcal{F}_i[k]$ s.t. $B \subset C \subseteq \Theta$. This is characteristic of an agent who is being cautious to incoming evidence due to high credibility of its existing knowledge. For instance, as we will discuss in the following section, an opinion leader, or a highly reliable estimate of the ground truth (GT), can be modeled as an agent employing a cCUE strategy who allows only refinement of the evidence it already has. Unlike previous work on consensus analysis, the work in this paper can accommodate a cautious agent thus allowing us to establish the relationship between the consensus and the GT.

4) Probabilistic Case: A probabilistic assignment $P_i[k] : \Theta \mapsto [0, 1]$ which is consistent with the evidence in $\mathcal{E}_i[k]$ satisfy $\mathrm{Bl}_i(B)[k] \leq P_i(B)[k] \leq \mathrm{Pl}_i(B)[k]$, $\forall B \subseteq \Theta$. Therefore, for probabilistic modeling (which only allows support to be assigned to singletons of Θ, and not composite), we have $P_i(B)[k] = m_i(B)[k] = \mathrm{Bl}_i(B)[k] = \mathrm{Pl}_i(B)[k]$. Thus, DST state modeling encapsulates probabilistic state models for which the CUE-based updates reduce to $P_i(B)[k+1] = \alpha_i[k]P_i(B)[k] + \sum_{j\neq i} C_{ij}[k]P_{ij}(B)[k]$, for rCUE and $P_i(B)[k+1] = P_i(B)[k]$, for cCUE.

Remark: Note that rCUE reduces to a weighted average of p.m.f.s. Therefore, rCUE belief revision in conjunction with DST state models encapsulates the traditional weighted average modeling approach of consensus. However, with the cCUE, the agent state remains the same. This in fact makes perfect sense—cCUE only allows refinements to originally cast evidence, and no further refinements are possible with probabilistic modeling, since only singleton propositions possess non-zero support.

5) Computational Complexity: A comparative quantitative analysis of the complexity of our algorithm with existing algorithms is difficult because of the different agent state model (viz., DST BoE) that we use. When the agent state (a DST BoE or a p.m.f.) is cast as a real-valued vector, the computational complexity/requirements related to the exchange of information between agents of our algorithm is comparable to those algorithms that employ a real-valued vector for their agent states [4]–[8], [10]–[12], [14], [16]–[18], except of course the major difference that we previously alluded to: the interaction protocols that we can use are severely restricted because of the restrictions placed on the agent state entries (viz., each entry must be non-negative and all entries must sum to one). This fundamental difference in our agent state model makes a reasonable comparison with existing works difficult. It would require a carefully constructed analytical treatment and an appropriate set of experiments which unfortunately could not be undertaken within the page limitations.

The generality of our agent state model (viz., DST BoE) offers an approach to study evidence updating and consensus dynamics in complex fusion environments where simpler models may not suffice. But this capability comes at a computational cost because a general DST BoE has $2^n - 1$ independent mass assignments. The recent work in [25] (and the references therein) offers a promising avenue to mitigate the associated complexity. Another approach is to utilize simpler agent models.

a) Dirichlet BoEs.: Dirichlet BoEs allow one to capture DST's ambiguity modeling capability with only a modest increase in cost. In standard probability, a p.m.f. has $n-1$ independent probability assignments to the focal elements $\{\theta_1, \ldots, \theta_n\}$. A Dirichlet BoE augments these singleton focal elements with only one additional focal element Θ which allows one to capture complete ambiguity. When all the BoEs are Dirichlet, the CUE in Definition 10 reduces to yield

$$m_i(B)[k+1] = \begin{cases} \alpha_i[k] m_i(B)[k] \\ \quad + \sum\limits_{j \neq i} \left[\beta_{ij}(B)[k] + \beta_{ij}(\Theta)[k] m_j(B)[k]\right], \\ \quad \text{for } B = \theta_\ell, \ \ell = 1, \ldots, n; \\ \alpha_i[k] m_i(B)[k] + \sum\limits_{j \neq i} \beta_{ij}(\Theta)[k] m_j(\Theta)[k], \\ \quad \text{for } B = \Theta, \end{cases} \tag{2}$$

i.e., the updated BoE remains Dirichlet. Comparing this with the probabilistic case (see Section III-C-4), it is clear that the computational complexity of updating the Dirichlet BoE case with n singletons is comparable to the computational complexity of updating the probabilistic case with $n+1$ singletons.

b) Pignistic P.M.F.s.: At any stage of belief revision, one can convert a DST BoE to its pignistic p.m.f. generated as

$$P(\theta_i) = \sum_{\theta_i \in A \subseteq \Theta} \frac{m(A)}{|A|}, \tag{3}$$

which possesses the property that $\mathrm{Bl}(A) \leq P(A) \leq \mathrm{Pl}(A), \ \forall A \subseteq \Theta$ [37]. Of course, reverting back to a p.m.f. is not reversible and evidence contained in non-singleton propositions is completely lost.

D. Modeling of Agent Interactions

Both *spatial* and *temporal* coupling among agents play key roles in the convergence characteristics of a multi-agent system.

1) Spatial Coupling: Spatial coupling among agents is best described via a graph. Consider the state update that takes place at time t_k via $F^{I[k]}$: the agent $\mathcal{A}_i$, where $i \leftarrow I[k]$, updates its state by interacting with agents $\mathcal{A}_j \in \mathcal{N}_{I[k]}$. The rooted digraph $G[k] = (V[k], E[k])$, with the vertex set $V[k] = \{\mathcal{A}_i \cup \mathcal{N}_{I[k]}\}$ and the edge set $E[k] = \{e_{ji} | \mathcal{A}_j \in \mathcal{N}_{I[k]}\}$, represents unidirectional information flow from $\mathcal{A}_j \in \mathcal{N}_{I[k]}$ to $\mathcal{A}_i$ via the directed edge e_{ji}. Such a graph is referred to as an *agent interaction topology (AIT)* in [31]. The multi-agent system is said to be *fully-connected* if $V[k] = \mathcal{N}, \ \forall k \in \mathbb{N}$ (i.e., all AITs used throughout the belief revision process are fully-connected); otherwise the system is said to be *partially-connected*. Now, consider the set of all AITs used by agent $\mathcal{A}_i \in \mathcal{N}$, i.e., the set of graphs $\mathbb{G}_i = \{G[k] = (\{\mathcal{A}_i \cup \mathcal{N}_{I[k]}\}, E[k]) | i \leftarrow I[k]\}$. We say that the multi-agent system is *static*, if $|\mathbb{G}_i| = 1, \ \forall \mathcal{A}_i \in \mathcal{N}$ (i.e., each agent $\mathcal{A}_i$ uses a single AIT throughout the updating process); otherwise, it is said to be *dynamic*.

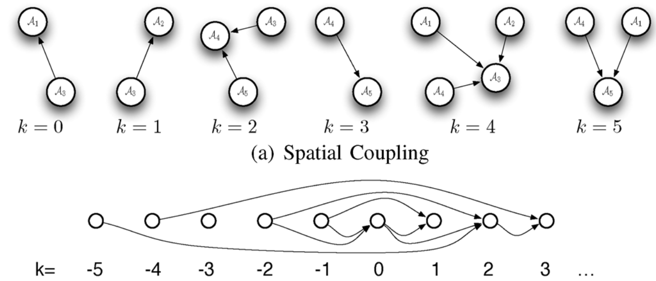

(a) Spatial Coupling

(b) Temporal Coupling

Fig. 2. **Spatial and temporal coupling of $\mathcal{N} = \{\mathcal{A}_1, \ldots, \mathcal{A}_5\}$ as given in Ex. 2**. Note that (a) spatial coupling specified via $F \in \mathcal{F}$ can only capture the agents that are involved in a given iteration. In an asynchronous setup, agent coupling alone cannot guarantee convergence. (a) Spatial coupling; (b) temporal coupling.

2) Temporal Coupling: When the communication among agents is asynchronous (e.g., when delays exists), spatial coupling alone is insufficient for convergence analysis. As illustrated in [27], the graph of an asynchronous iteration (see Definition 8) can be used to analyze temporal coupling among fusion operators. Fig. 2 illustrate the spatial and temporal coupling (of the corresponding asynchronous iteration; see Section V-A for details) of agents in Example 2.

E. Social Networks: A Case Study

Convergence characteristics of information diffusion in a social network can be analyzed from a data fusion perspective, where the agents, together with their interactions and communications, in a multi-agent system form a complex fusion environment. Preliminary work on the impact of temporal connectivity or *temporal dynamics* of communication on information flow in social networks is explored in [38]. *Underlying graph* or the spatial connectivity among agents in social networks have been discovered to possess special patterns that characterizes such networks. Some of the widely studied network topologies include the following:

(a) **Erdös-Rényi networks** are mainly characterized by randomly formed agent connectivity, thus having no systematic patterns of interactions (e.g., a new group of friends). The *edge density* $\gamma \in [0, 1]$ quantifies the connectivity (with $\gamma = 0$ for no connectivity, and $\gamma = 1$ for a fully-connected graph), where $\gamma < 1/m$ (for a graph with m vertices) characterizes a graph that consists of small groups and isolated individuals [39].

(b) **Small-world networks** characterize social interactions that are mostly local, but connected with some "long-range" interactions [40].

(c) **Scale-free networks** follow an inverse power law in the number of connections per vertex. Such networks are often formed when new agents "preferentially" attach themselves to agents who are highly connected to other agents. This is a common network topology for large social networks, world-wide-web (WWW) and communication networks [41].

(d) **Hierarchical networks or trees** are characteristic of social interactions such as family trees, organization structures, etc., where there are edges between every vertex pair such that no cycles or loops are formed [42].

IV. Paracontractions View of Consensus

With the notions in Section II-C in place, one can view a consensus state $\mathbf{x}^* \in \mathbb{D}$ in a multi-agent system as a common fixed point of the pool of belief revision operators used by agents. As we show in Section V, convergence of belief revision using a pool of (para)contractive operators that contains common fixed points can then be established using convergence results on asynchronous iterations. However, from a data fusion perspective, it is also important to adequately understand the characteristics of belief revision operators that yield a consensus state (if it is attained). We now derive a pool of CUE-based belief revision operators that are guaranteed to generate a *rational consensus*, and establish their contractive properties. Let us proceed by formally defining consensus.

A. Consensus Notions

1) Consensus in a Fusion Environment: The word consensus refers to a general agreement [43] among sources (e.g., a consensus of opinions among a jury pool), and has been adopted in many sensor research/applications [4], [6], [9], [10], [11], [14], [31] where an "agreement" among sensor measurements is sought. Consensus, as applicable to a fusion environment, is taken as follows.

Definition 11 (Consensus): Consider the multi-agent system $\mathcal{N} = \{\mathcal{A}_1, \ldots, \mathcal{A}_m\}$ of adaptive agents that interact towards collectively estimating $\mathcal{X}$. The initial conditions are taken as $\mathcal{X}_0 = \{\mathcal{E}_i[0] | i = 1, \ldots, m\}$. Each agent $\mathcal{A}_i$ iteratively updates its state $\mathcal{E}_i[k] = \{\Theta, \mathcal{F}_i[k], m_i(\cdot)[k]\}$ at t_k via some fusion operator $F \in \mathcal{F}_\triangleleft$ as identified by the sequences $\mathcal{I}$ and $\mathcal{S}$. Then, we say 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 $\|\cdot\| : \mathcal{E}_\Theta \mapsto \Re_{+0}$. □

2) Rational Consensus: The seminal work in [1], which has since become popular as the *Lehrer-Wagner model of rational consensus*, espouses the weighted average as the basis for consensus formation. This constitutes the core of many consensus related papers. However, in complex fusion environments that are rife with data uncertainties, satisfying this requirement alone is not sufficient towards generating a rational consensus. One reassuring property in a majority of soft/hard fusion environments is the availability of ground truth GT estimates. These estimates $\widehat{GT}$ are often vague, but highly reliable (e.g., prevailing threat level estimates extracted via satellite imagery, initial hypotheses of opinion leaders in a social network setting). Hence, we believe a more meaningful, or 'rational' consensus must satisfy the following properties: $\boldsymbol{\mathcal{P}_1}$: it is based on a weighted average of agent states [1]; $\boldsymbol{\mathcal{P}_2}$: it is consistent with $\widehat{GT}$ when it is available; and $\boldsymbol{\mathcal{P}_3}$: it converges to the GT when it is known. More formally, in terms of DST BoEs, these properties of a rational consensus can be expressed as

Definition 12 (Rational Consensus): Let $\mathcal{E}^t = \{\Theta, \mathcal{F}^t = \theta^t \in \Theta, m^t(\theta^t) = 1.0\}$ and $\widehat{\mathcal{E}}^t = \{\Theta, \widehat{\mathcal{F}}^t, \widehat{m}^t(\cdot)\}$ denote the GT and $\widehat{GT}$, respectively. Let the BoE $\mathcal{E}^* = \{\Theta, \mathcal{F}^*, m^*(\cdot)\}$ denote a consensus state reached by the agents in $\mathcal{N}$ via some belief revision process. We say $\mathcal{E}^*$ is **rational** if

(i) $\mathcal{E}^*$ is reached as some convergence point of weighted averages of $\mathcal{E}_i[\cdot], i = 1, \ldots, m$;

(ii) $\mathcal{E}^*$ is a refinement of $\widehat{\mathcal{E}}^t$ when it is known (i.e., for all $B \in \mathcal{F}^*$, there is a $C \in \widehat{\mathcal{F}}^t$ s.t. $B \subseteq C$); and

(iii) $\mathcal{E}^* \to \mathcal{E}^t$, whenever $\widehat{\mathcal{E}}^t = \mathcal{E}^t$ is known. □

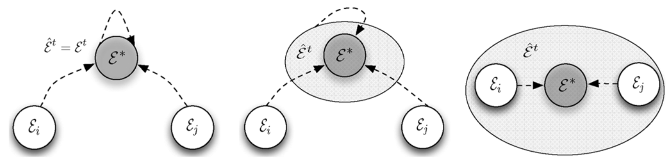

(a) GT is known (b) $\widehat{GT}$ is known (c) Unknown GT

Fig. 3. **Graphical illustration of rational consensus**. Clearly, $\mathcal{E}^t$ must take the form s.t. $m^t(\theta^t) = 1.0$ for some $\theta^t \in \Theta$, where θ^t is the GT. Since $\widehat{\mathcal{E}}^t$ is an estimate of the GT, we assume that $\exists B \in \mathcal{F}^t$ s.t. $\theta^t \in B$. (a) GT is known; (b) $\widehat{GT}$ is known; (c) Unknown GT.

Fig. 3 illustrates the nature of a rational consensus as defined in Definition 12.

As we will show presently, it turns out that a particular CUE parameterization (see Section II-C-2) guarantees the formation of a rational consensus as defined in Definition 12. We will also show that, under certain mild conditions, this parameterization is also contractive on $\mathcal{E}_\Theta$, which is a necessary condition for convergence analysis under the umbrella of asynchronous iterations (see Theorems 3 and 4).

B. CUE-Based Paracontracting Pool of Operators

Consider the following CUE-based fusion operator.

Definition 13 ($F_\triangleleft$ Operator): Consider a set of BoEs $\mathcal{E}_i \in \mathcal{E}_\Theta, i = 1, \ldots, m'$, that possibly contains $\hat{\mathcal{E}}^t$ (i.e., $\exists j \in \{1, \ldots, m'\}$, s.t. $\mathcal{E}_j \equiv \hat{\mathcal{E}}^t$, whenever $\hat{\mathcal{E}}^t$ exists). Then, the operator $\mathrm{F}^i_\triangleleft : \mathcal{E}_\Theta^{m'} \mapsto \mathcal{E}_\Theta$ that updates $\mathcal{E}_i$ with the evidence in $\mathcal{E}_j, j \in \{1, \ldots, m'\} \setminus \{i\}$ is defined via the CUE as $\mathrm{F}^i_\triangleleft(\mathcal{E}_1, \ldots, \mathcal{E}_{m'}) = \mathcal{E}_i \triangleleft_{j=1}^{m'} \mathcal{E}_j$, where the CUE parameters are given by

$$\alpha_i = \lambda_i \neq 1; \quad \beta_{ij}(A) = \begin{cases} \lambda_j m_i(A), & for\ \mathcal{E}_i = \hat{\mathcal{E}}^t; \\ \lambda_j m_j(A), & otherwise, \end{cases}$$

s.t. $\alpha_i + \sum_{j \neq i} \sum_{A \in \mathcal{F}_j} \beta_{ij}(A) = 1$, for $\lambda_i \in \Re_{+0}, i = 1, \ldots, m'$. □

Note that the parameter choices are based on rCUE and cCUE strategies: cCUE is used for updating the $\widehat{GT}$, thus only allowing refinements on initial assignments; rCUE is used for the other agents, so that they are receptive to information coming from other agents. $F_\triangleleft$ definition allows the inclusion of multiple $\hat{\mathcal{E}}^t$ in the mix as long as they are "consistent" (or do not contradict) one another. We will address the exact requirements for inclusion of multiple $\widehat{GT}$ after establishing the following results.

Claim 5: The $F_\triangleleft$ operator in Definition 13 has infinitely many fixed points in $\mathcal{E}_\Theta$. Furthermore, if $\mathcal{E}^* \in \mathrm{fix}(\mathrm{F}_\triangleleft)$, then for all $B \in \mathcal{F}^*$, $\nexists C \in \mathcal{F}^*$ s.t. $B \subset C$ or $B \supset C$. □

Claim 5 establishes the fact that the operator $\mathrm{F}_\triangleleft$ can in fact generate a consensus, viz., some fixed-point in $\mathcal{E}_\Theta$. Furthermore, it also states that the consensus BoE will have a unique structure: any proposition (with non-zero support) will not *contain* or be *contained in* any other proposition (see Fig. 4 for a set theoretic interpretation). Therefore, consensus BoEs will only contain propositions that do not allow further refinements. However, such propositions need not be singletons.

Remark: Note, that this scenario is not important in pure probabilistic assignments, since all propositions with non-zero

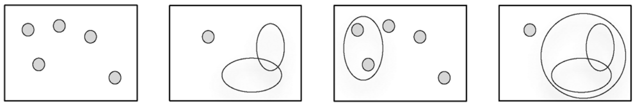

(a) $\mathfrak{F}_1^*$ (b) $\mathfrak{F}_2^*$ (c) $\mathfrak{F}_1$ (d) $\mathfrak{F}_2$

Fig. 4. **The structure of consensus BoEs**. Shaded and non-shaded objects represent singleton and composite propositions, respectively. $\mathcal{F}_i^*$, $i = 1, 2$, represent two examples of fixed points of $F_\triangleleft$; $\mathcal{F}_i$, $i = 1, 2$, are example that cannot be fixed points of $F_\triangleleft$. (a) $\mathcal{F}_1^*$; (b) $\mathcal{F}_2^*$; (c) $\mathcal{F}_1$; (d) $\mathcal{F}_2$.

support are elements of Θ (i.e., $\forall B \in \mathcal{F} \implies B \in \Theta$). However, in the case of partial probability specifications (as allowed in DST), BoEs are allowed to contain composite propositions (i.e., $B \subseteq \Theta$) to represent uncertainty. Claim 5 establishes that, even under this case, consensus can only contain propositions that cannot be further refined. The assurance that the generated consensus state will be the 'most' refined as allowed by evidence is in fact a useful observation especially in the case of soft/hard fusion environments, where soft sources often provide vague evidence (hence composite propositions).

C. Reaching a Consensus

With Claim 5 in place, we are now in a position to define a DST *consensus protocol* (i.e., a scheme that can generate a consensus via iterative belief revision) for complex fusion environments.

Definition 14: Let $\mathcal{N} \equiv \{\mathcal{A}_1, \ldots, \mathcal{A}_m\}$ be a multi-agent system, where agent states and the initial conditions are given via $\mathcal{E}_i[k] = \{\Theta, \mathcal{F}_i[k], m_i(\cdot)[k]\}, i = 1, \ldots, m$, and $\mathcal{X}_0 = \{\mathcal{E}_i[0] | i = 1, \ldots, m\}$, respectively. Then, for an operator sequence $\mathcal{I} \equiv (I[k] \in \mathbb{I})_{k=0}^{\infty}$, a DST asynchronous consensus protocol on $\mathcal{N}$ is given via

$$\mathcal{E}_j[k+1] = \begin{cases} \mathrm{F}_\triangleleft^{I[k]}\left(\mathcal{E}_{i_1}\left[s^1(k)\right], \ldots, \mathcal{E}_{i_{m_{I[k]}}}\left[s^{m_{I[k]}}(k)\right]\right), & \mathit{for}\ j = i \leftarrow I[k]; \\ \mathcal{E}_j[k], & \mathit{for}\ j \neq i. \end{cases}$$

Here, $\mathcal{S} = \{S[k]\}_{k=0}^{\infty}$ denotes a sequence of m_i-tuples $S[k] = (s^1[k], \ldots, s^{m_{I[k]}}[k])$ with $s^\ell[k] \in \mathbb{N}_0 \cup \{-1, \ldots, -m\}$ s.t. $s^\ell[k] \leq k$, $\forall \ell \in \{1, \ldots, m_{I[k]}\}$, where $m_i \leq m$, $\forall i \in \mathbb{I}$. $\square$

Remark: All agents $\mathcal{A}_i \in \mathcal{N}$ start off with the initial states $\mathcal{E}_i[0]$ and update their states in the sequence specified via $\mathcal{I}$. The state updating process will take the following form:

(i) The BoE of an agent which is not a $\widehat{GT}$ will receive all the focal elements from its neighbors, i.e., $\mathcal{F}_i[k+1]$ of an agent $\mathcal{A}_i, i \leftarrow I[k]$, via neighbor set $\mathcal{N}_{I[k]}$ will have $\mathcal{F}_i[k+1] = \mathcal{F}_i[k] \cup_{j \in \mathcal{N}_{I[k]}} \mathcal{F}_j[k]$;

(ii) The BoE of an agent having a $\widehat{GT}$ will only receive "new" focal elements that are "refinements" to its existing focal set, i.e., $B \in \mathcal{F}_i[k+1] \setminus \mathcal{F}_i[k]$ iff $\exists C \in \mathcal{F}_i[k]$ s.t. $B \subseteq C$. After sufficient number of initial iterations, (**ii.a**) all agents who do not possess a $\widehat{GT}$ will have received focal elements of all other agents directly or indirectly (i.e., via another agent who has updated its state with other agents); and (**ii.b**) all agents who possess a $\widehat{GT}$ will have received all focal elements from all other agents that are a refinement to initial state.

We now show that the iteration operators $\mathrm{F}_\triangleleft^i, i \in \mathbb{I}$, of the form in Definition 14 are paracontractive on $\mathcal{E}_\Theta$. This sets the foundation for CUE-based consensus analysis.

Claim 6: When used in an iteration of the form Definition 14, the $\mathrm{F}_\triangleleft$ operator in Definition 13 is paracontractive on $\mathcal{E}_\Theta$ w.r.t. the p-norm $\|\mathcal{E}\|$ where $\|\mathcal{E}\|^p = \sum_{B \subseteq \Theta} |m(B)|^p$. $\square$

Another result which is necessary for reaching a consensus state under the iterations given in Definition 14 is

Claim 7: Let $\mathcal{F}_\triangleleft \equiv \{F_\triangleleft^i : \mathcal{E}_\Theta^{m_i} \mapsto \mathcal{E}_\Theta | m \geq m_i \in \mathbb{N}, i \in \mathbb{I}\}$ denote the pool of $F_\triangleleft$ operators enumerated via the index set $\mathbb{I}$ for updating agents in $\mathcal{N}$ with initial conditions $\mathcal{X}_0 \equiv \{\mathcal{E}_i[0] | i = 1, \ldots, m\}$. Then, the pool $\mathcal{F}_\triangleleft$ contains common-fixed points under following conditions.

(i) $\mathcal{F}_\triangleleft$ only contains operators corresponding to at most one agent having access to a $\widehat{GT}$.

(ii) If $\mathcal{F}_\triangleleft$ contains operators corresponding to more than one agent having access to $\widehat{GT}$; then, the initial conditions $\mathcal{E}_i[0], \mathcal{E}_j[0] \in \mathcal{X}_0$ of such agents are s.t. $\exists C \in \mathcal{F}_j^t[0]$ for which $B \subseteq C$ or $B \supseteq C$, for all $B \in \mathcal{F}_i^t[0]$.

$\square$

Proof: Part (***i***) immediately follows from Claim 5, since any fixed point of operators corresponding to $\widehat{GT}$ of the form given in Claim 5 are fixed-points of all other operators.

For part (***ii***) notice that BoE of an agent having a $\widehat{GT}$ will only receive "new" focal elements that are "refinements" to its existing focal set, i.e., $B \in \mathcal{F}_i[k+1] \setminus \mathcal{F}_i[k]$ iff $\exists C \in \mathcal{F}_i[k]$ s.t. $B \subseteq C$. Now, given all BoEs corresponding to $\widehat{GT}$ have initial conditions $\mathcal{E}_i[0], \mathcal{E}_j[0] \in \mathcal{X}_0$ are s.t. $\exists C \in \mathcal{F}_j^t[0]$ for which $B \subseteq C$ or $B \supseteq C$ for all $B \in \mathcal{F}_i^t[0]$, any fixed-point $\mathcal{E}^*$ of the form given in Caim 5 s.t. all $B \in \mathcal{F}^*$ is *contained in* some $C \subseteq \Theta$ s.t. common to $\mathcal{F}_j^t[0]$ for all initial conditions $\mathcal{E}_j[0]$ corresponding to $\widehat{GT}$. Then, the claim follows since such fixed-points are also common to all other operators in $\mathcal{F}_\triangleleft$. $\blacksquare$

This claim establishes the conditions which leads to consensus state generation via the iteration in Definition 14 using the pool $\mathcal{F}_\triangleleft$ of fusion operators. Hereinafter, we assume that the initial conditions $\mathcal{X}_0$ satisfies the conditions regarding multiple $\widehat{GT}$.

The belief revision process given in Definition 14 employs fusion operators from $\mathcal{F}_\triangleleft$. Thus the Claims 5, 6, and 7 on fixed-points and contractive properties in $\mathcal{E}_\Theta$ are satisfied. In fact, we now have

Claim 8: A consensus BoE generated via the protocol in Definition 14 is a rational consensus (Definition 12). $\square$

Proof: Claim is easily established by utilizing CCT and fusion properties of CUE. $\blacksquare$

Claim 8 guarantees the ability of the iteration in Definition 14 to generate a rational consensus in complex fusion environments. However, convergence of belief revision in such environments is governed by the properties of fusion operators as well as the characteristics of agent interactions (i.e., $\mathcal{I}$ and $\mathcal{S}$). Let us now study the conditions on the sequences $\mathcal{I}$ and $\mathcal{S}$ that will guarantee convergence.

V. Convergence Analysis of a CUE-Based Pool

Now that we have established the ability of CUE-based fusion operators to generate a rational consensus, in this section, we make use of Theorems 3 and 4 to explore the criteria (viz., the conditions on $\mathcal{I}$ and $\mathcal{S}$) that will guarantee convergence of CUE-based belief revision under several network topologies.

A. Belief Revision Process as an Asynchronous Iteration

In an asynchronous iteration as given in Definition 6, at index k, $\mathbf{x}[k] \in \mathbb{D}$ is updated via an operator $F^{I[k]} \in \mathcal{F}_{\triangleleft}$. However, in a consensus setup, at index k, state of agent $\mathbf{x}_i[k] \in \mathbb{D}, i \leftarrow I[k]$, is updated via $F^{I[k]}$. However, the equivalence of these two iterations can be easily established as shown in [31]. The idea behind this simple transformation is to renumber the sequences $\mathcal{I}$ and $\mathcal{S}$ in such a manner that the vector $\mathbf{x}[k]$ being updated at k via $F^{I[k]}$ is equivalent to updating $\mathbf{x}_i[k']$ via $F^{I[k']}$ at some index k', where $t_k = t'_k$ (note that, having $k \neq k'$ is irrelevant as long as $t_k = t_{k'}$). We provide the relevant transformation as given in [31].

Step 1) Renumber $s^\ell[k], k = 0, 1, \ldots$, in such a manner that all the components $\mathbf{x}_\ell[s^\ell[k]]$ in Definition 9 are updated at time $s^\ell[k]$, i.e.,

$$I\left(s^i[k]-1\right) = i, \quad s^i[k] \geq 1, \quad i \in \{1,\ldots,m\}, \quad \forall k \in \mathbb{N}. \tag{4}$$

Step 2) Write all initial conditions as vector multiples of $\mathbf{1}$ (identity matrix) and set $\mathbf{x}[-k] = \mathbf{x}_k[0]\mathbf{1}$, $\forall k = 1,\ldots,m$, and renumber the elements for which $s^\ell[k] = 0$ in the same way for $\ell \in \{1,\ldots,m\}$ and $k \in \mathbb{N}_0$.

Step 3) Generate the asynchronous iteration $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \tilde{\mathcal{S}})$ with

$$\mathbf{y}[k+1] = F_{\triangleleft}^{I[k]}\left(\mathbf{y}\left[\tilde{s}^1[k]\right], \ldots, \mathbf{y}\left[\tilde{s}^{m_{I[k]}}[k]\right]\right), \quad k \in \mathbb{N}_0, \tag{5}$$

where $\tilde{\mathcal{S}}$ is a sequence of m_i-tuples with elements given by

$$\tilde{s}^i[k] = s^{m_{I[k]}}[k], \quad i = 1,\ldots,m_{I[k]}, \quad \forall k \in \mathbb{N}_0, \tag{6}$$

and $\mathcal{Y}_{\mathcal{O}}$ consists of $\mathbf{y}[-\ell] = \mathbf{x}[-\ell], \ell = 1,\ldots,m$.

Now, the asynchronous iteration $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \tilde{\mathcal{S}})$ is equivalent to the belief revision process in Definition 9 in the sense that

$$\mathbf{y}[k+1] = \mathbf{x}_i[k+1], \quad i \leftarrow I[k], \quad \forall k \in \mathbb{N}_0. \tag{7}$$

B. Verification of Convergence

We now establish the conditions that allow us to use Theorem 3. A result in [31] that can be used to simplify the conditions for a confluent iteration in Definition 8 is

Lemma 9: [31] Let $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \tilde{\mathcal{S}})$ be the asynchronous iteration in (7). Consider any trajectory of $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \tilde{\mathcal{S}})$ along the sequence of AITs $G[0], G[1], \ldots$. Suppose this sequence $G[0], G[1], \ldots$ is repeatedly jointly rooted, and satisfies the following assumptions:

(i) Agent $\mathcal{A}_{i_0}$ always uses its own latest state to update its current state, i.e., $s^{i_0}(k) = \max\{k_0 \leq k | I[k_0-1] = i_0\}$ for all $k > \min\{k_0 \in \mathbb{N}_0 | I[k_0] = i_0\}$ with $I[k] = i_0$;
(ii) $\mathcal{I}$ is regulated; and
(iii) $k - s^\ell[k] \leq s$, $\forall k \in \mathbb{N}_0, \ell = 1,\ldots,m$, for some $s \in \mathbb{N}_0$.

Then, $(\mathcal{F}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \tilde{\mathcal{S}})$ is confluent. □

Remarks: (**a**) Theorem 3 states that, if $\mathcal{F}_{\triangleleft}$ is a finite paracontracting pool with common fixed-points, then a confluent asynchronous iteration $(\mathcal{F}_{\triangleleft}, \mathcal{X}_0, \mathcal{I}, \mathcal{S})$ converges to some common fixed point of $\mathcal{F}_{\triangleleft}$. We have already established the existence of common fixed points and paracontracting properties of the CUE-based pool via Claims 6 and 7. Therefore, if $\mathcal{F}_{\triangleleft}$ (or equivalently $\mathbb{I}$) is finite, then convergence criteria can be derived by deriving conditions that satisfy conditions for confluent iterations as given in Lemma 9.

(**b**) In the case of an infinite operator pool, one can make use of Theorem 4 for establishing convergence. For instance, in belief revision scenarios, it is common to use relative distance for 'weighing' the incoming evidence. In this case, the parameters $\lambda_{(\cdot)}$ corresponding to the update

$$\mathcal{E}_i[k+1] = \mathrm{F}_{\triangleleft}^{I[k]}\left(\mathcal{E}_{i_1}\left[s^1(k)\right], \ldots, \mathcal{E}_{i_{m_{I[k]}}}\left[s^{m_{I[k]}}(k)\right]\right) \tag{8}$$

will usually take the form $\lambda_{i_j} \propto \|\mathcal{E}_i[k] - \mathcal{E}_{i_j}[s^{i_j}(k)]\|$, thus potentially leading to an infinite pool of operators. However, given the continuity of CUE-based fusion (which directly follows from the continuity of DST belief functions), a finite pool that is approximated by such an infinite pool $\mathcal{G} = \{G^k | k \in \mathbb{N}_0\}$ can be easily derived (by showing that there exists a null sequence ϵ_j and a sequence $\mathcal{I}' = \{i_k \in \mathbb{I}\}_{k=0}^{\infty}$ s.t. $\|G^k(\cdot) - F^{i_k}(\cdot)\| \leq \epsilon_k, k \in \mathbb{N}_0$; see Definition 5). Now, if the initial conditions in such a pool are made to satisfy $\mathcal{D}_0$ as given in Theorem 4, we can again utilize the conditions on confluent iterations to study convergence. Therefore, without loss of generality, we focus our analysis to a finite pool $\mathcal{F}_{\triangleleft}$ of CUE-based operators that contains common fixed points.

Now we use the equivalence of iterated belief revision in a consensus setup and asynchronous iterations along with Lemma 9 for convergence analysis of CUE-based belief revision under several network configurations.

1) Synchronous, Fully-Connected Network: This represents perhaps the simplest multi-agent system, where each agent is connected to all the other agents and information is exchanged without any iteration delay (i.e., $k - s^\ell[k] = 0$, $k \in \mathbb{N}_0$). In this case, iterated belief revision in Definition 14 reduces to

$$\mathcal{E}_j[k+1] = \begin{cases} \mathrm{F}_{\triangleleft}^{I[k]}\left(\mathcal{E}_1[k], \ldots, \mathcal{E}_m[k]\right), & j = i \leftarrow I[k]; \\ \mathcal{E}_j[k], & j \neq i. \end{cases} \tag{9}$$

Claim 10: If $\mathcal{I}$ is regulated, the iterated belief revision in (9) converges. □

Proof: Given that $\mathcal{I}$ is regulated, we first prove that the equivalent iteration $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \mathcal{S})$ in (9) is confluent. Clearly, the sequence of AITs $\{G[k]\}_{k=1}^{\infty}$ is repeatedly jointly rooted. (**a**) Each agent $\mathcal{A}_i, i = 1,\ldots,m$, uses its latest state for the update; (**b**) $\mathbb{I} = \{1,\ldots,m\}$ is regulated; and (**c**) $k - s^\ell[k] = 0 \leq s$, $\forall k \in \mathbb{N}_0, \ell = 1,\ldots,m$, for all $s \in \mathbb{N}_0$.

So, Lemma 9 yields that $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \mathcal{S})$ is confluent. Now, according to Theorem 3, the iteration in (9) converges. ■

2) Synchronous, Static, Partially-Connected Network: Here, agents communicate without iteration delays (i.e., $k - s^j[k] = 0$) with (at least one) partially connected AIT that does not change over time (e.g., hierarchical, static scale-free, networks). Therefore, some agents cannot communicate with others. In this case, the consensus protocol in Definition 14 reduces to

$$\mathcal{E}_j[k+1] = \begin{cases} \mathrm{F}_{\triangleleft}^{I[k]}\left(\mathcal{E}_{i_1}[k], \ldots, \mathcal{E}_{i_{m_i}}[k]\right), & j = i \leftarrow I[k]; \\ \mathcal{E}_j[k], & j \neq i. \end{cases} \tag{10}$$

Claim 11: If $\mathcal{I}$ is regulated, iterated belief revision in (10) converges as long as the graph union of AITs of all agents is connected. □

Proof: Consider the equivalent iteration $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \mathcal{S})$ of (10). Since, graph union of AITs is connected and each agent uses its own latest state for updating, the sequence of AITs $G[k]$ are repeatedly jointly rooted. Similar to the case of a synchronous fully-connected network, clearly assumptions (i) and (iii) of Lemma 9 are satisfied. Hence, $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \mathcal{S})$ is confluent according to Lemma 9. Thus, the iteration in (10) converges. ■

3) Synchronous, Dynamic, Partially-Connected Network: Here, agents communicate without iteration delays (i.e., $k - s^j[k] = 0$), and with (at least one) partially connected AITs that change over time (e.g., Erdös-Rényi, ad-hoc, networks). This setup is similar to the above, except that the neighbor set of agents change over time. In this case, iterated belief revision in Definition 14 reduces to

$$\mathcal{E}_j[k+1] = \begin{cases} \mathrm{F}_{\triangleleft}^{I[k]}\left(\mathcal{E}_{i_1}[k], \ldots, \mathcal{E}_{i_{m_{I[k]}}}[k]\right), & j = i \leftarrow I[k]; \\ \mathcal{E}_j[k], & j \neq i, \end{cases} \tag{11}$$

where $m_{I[k]} \in \{1, \ldots, m\}$. As above, if $\mathcal{I}$ is regulated, the iteration in (11) will converge whenever the AITs $G[k], k \geq 1$, of the network are repeatedly jointly rooted.

4) Asynchronous, Fully-Connected Network: Here, each agent is connected to all the other agents, but the information exchange is asynchronous (or has delays) (i.e., $k - s^{\ell}[k] < 0$). In this case, iterated belief revision in Definition 14 reduces to

$$\mathcal{E}_j[k+1] = \begin{cases} \mathrm{F}_{\triangleleft}^{I[k]}\left(\mathcal{E}_1\left[s^1[k]\right], \ldots, \mathcal{E}_m\left[s^m[k]\right]\right), & j = i \leftarrow I[k]; \\ \mathcal{E}_j[k], & j \neq i. \end{cases} \tag{12}$$

Claim 12: If $\mathcal{I}$ is regulated, iterated belief revision in (12) converges whenever the iteration delays $k - s^{\ell}[k]$ are finite. □

Proof: Consider the equivalent iteration $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \mathcal{S})$ of (10). Since the network is fully-connected, the sequence of AITs $G[k], k \geq 1$, are repeatedly jointly rooted. Clearly, $(\mathcal{F}_{\triangleleft}, \mathcal{Y}_{\mathcal{O}}, \mathcal{I}, \mathcal{S})$ satisfies the assumptions (i) and (ii) of Lemma 9. Thus, if the iteration delays given by $k - s^{\ell}(k) \leq \infty \; \forall k \in \mathbb{N}_0, \ell = 1, \ldots, m$, are finite, then according to Lemma 9, belief revision in (12) converges.

Remark: Similar results can be obtained for asynchronous, dynamic, and partially-connected networks. In this case, similar to the cases we have studied above, one needs to impose conditions to guarantee adequate coupling among agents to reach a consensus. Given that the pool $\mathcal{F}_{\triangleleft}$ of CUE-based operators is paracontractive on $\mathcal{E}_{\Theta}$ (under the conditions on $\widehat{GT}$), the belief revision given in Definition 14 is capable of generating a rational consensus in any network setup as long as these coupling conditions are satisfied. In the analysis of an arbitrary network topology, one needs to make sure that it satisfies Lemma 9; and, then the convergence is automatically guaranteed via Theorem 3.

VI. Experiments

Here we study the convergence behavior of the proposed CUE-based belief revision under several network topologies, and highlight several key points as we have discussed earlier. All experiments were carried out on Matlab (version 2013b, licensed to University of Notre Dame du lac, Notre Dame, IN). The codes were written to simulate the iteration process without approximations or optimizations.

A. Setup

We generate a multi-agent system that consists of 100 agents, where each agent has an initial belief on the phenomenon $\mathcal{X}$ whose truth lies in $\Theta = \{a, b, c, d, e\}$.

1) Initially Cast Evidence: The initial conditions of the multi-agent system given by $\mathcal{E}_i[0]$, $i = 1, \ldots, 100$, are generated as 'noisy' versions of $\tilde{\mathcal{E}}_i = \{\Theta, \tilde{\mathcal{F}}_i, \tilde{m}_i(\cdot)\}, i = 1, \ldots, 5$, which represent five (5) imperfect sources of evidence whose masses are given by:

BPA	a	c	d	e	ac	ad	abc	bc	cd	cde
$\tilde{m}_1$	0.6	-	0.2	-	-	-	-	0.2	-	-
$\tilde{m}_2$	-	0.8	-	-	-	-	-	-	0.2	-
$\tilde{m}_3$	-	-	0.2	-	-	-	0.8	-	-	-
$\tilde{m}_4$	-	-	-	0.2	0.8	-	-	-	-	
$\tilde{m}_5$	-	-	-	-	-	0.8	-	-	-	0.2

Then, $\mathcal{E}_i[0], i = 1, \ldots, 100$, is generated as follows: (**a**) $\mathcal{F}_i[0] = \tilde{\mathcal{F}}_{\lfloor (i-1)/20 \rfloor + 1}$; (**b**) $m_i(B)[0] = \tilde{m}_{\lfloor (i-1)/20 \rfloor + 1}(B) + w_i, \; \forall B \in \tilde{\mathcal{F}}_{\lfloor (i-1)/20 \rfloor + 1}$, where w_i is Gaussian distributed with 0 mean and 0.1 variance; and finally, (**c**) renormalize $m_i(\cdot)[0]$ s.t. $\sum_{B \subseteq \Theta} m_i(B)[0] = 1$. These initial conditions mimic a scenario where agents have received information from several sources (5 in this case) and have developed their own subjective beliefs.

2) Agent Interactions: We employ 8 spatial configurations (both static and dynamic) and convergence is studied under both synchronous (i.e., without delays) and asynchronous (i.e., with delays) updates. The resulting configurations are:

Configuration			
Synchronous	**Static/Dynamic**	**Spatial Connectivity**	**Asynchronous**
C1	static	fully-connected	**C1d**
C2	static	partially-connected	**C2d**
C3	static	small-world	**C3d**
C4	static	scale-free	**C4d**
C5	dynamic	partially-connected	**C5d**
C6	dynamic	Erdös-Rényi	**C6d**
C7	dynamic	small-world	**C7d**
C8	dynamic	scale-free	**C8d**

Here, agent interactions corresponding to scale-free, small-world and Erdös-Rényi topologies were created using network generators simulating the link formations mechanisms as given in [39], [40], [44], respectively. For each topology, dynamic agent interactions were generated by sampling from a pool of 10 static networks. Asynchronous communications were simulated from the same spatial configurations, where the state updates were delayed according to a delay sequence $\mathcal{S}$ that was generated as an accumulation of uniformly distributed pseudorandom numbers generated at each iteration, i.e., $s^{(\cdot)}[k] = k - \sum_{j=0}^{k} d[k]$, where $d[k] \sim U[0, 1]$.

B. Characteristics of CUE-Based Iterations

We now study several interesting features of the CUE-based iterated belief revision. Except for Fig. 8, we only show the results for configuration C1 (i.e., fully-connected, static and no delays). In all the plots, the x-axis denotes the iteration k.

1) Convergence to $\widehat{GT}$: A very important feature of the proposed CUE-based consensus protocol is its ability to drive the consensus state towards a rational consensus (in the sense of

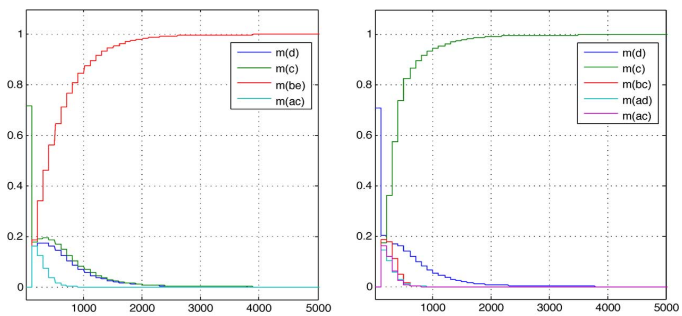

(a) $\widehat{GT} = \{be\}$ (b) $\widehat{GT} = \{bc\}$

Fig. 5. **Convergence of state of $\mathcal{A}_2$ under different $\widehat{GT}$**. Mass $m_2(\cdot)[k]$ appears on the y-axis. Initially cast evidence for $\mathcal{A}_2$ supports the two propositions $\{c\}$ and $\{cd\}$ only. (**a**) When $\widehat{GT} = \{be\}$: the consensus does not further refine $\widehat{GT}$. (**b**) When $\widehat{GT} = \{bc\}$: the consensus refines $\widehat{GT}$ and generates $\{c\}$ as the consensus. (a) $\widehat{GT} = \{be\}$; (b) $\widehat{GT} = \{bc\}$.

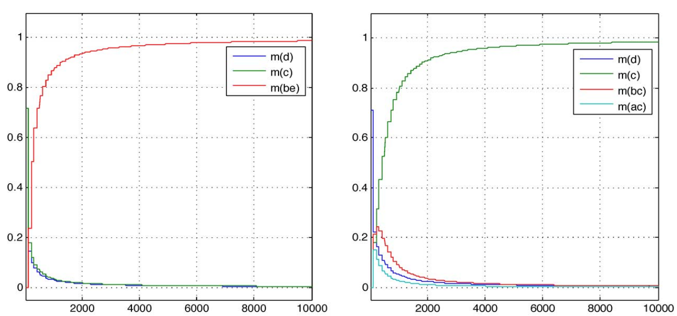

(a) $\widehat{GT} = \{be\}$ (b) $\widehat{GT} = \{bc\}$

Fig. 6. **Convergence behavior of $\mathcal{A}_2$ under an infinite pool of CUE operators**. Mass $m_2(\cdot)[k]$ appears on the y-axis. Note the initially cast evidence for $\mathcal{A}_2$ only support the two propositions $\{c\}$ and $\{cd\}$. While consensus states are similar to Fig. 5, convergence rate is slower. (a) $\widehat{GT} = \{be\}$; (b) $\widehat{GT} = \{bc\}$.

Definition 12). So, the consensus is guaranteed to be a refinement of a $\widehat{GT}$, when such information is available. Fig. 5 shows the convergence of agent $\mathcal{A}_2$ under two cases: (**a**) $\widehat{GT} = \{be\}$: The initially cast evidence of agent $\mathcal{A}_2$ does not support either proposition $\{b\}$ or $\{e\}$. Therefore, the consensus offers no further refinement of $\widehat{GT} = \{be\}$. (**b**) $\widehat{GT} = \{bc\}$: In this case, the initial beliefs of $\mathcal{A}_2$ are further refined, as the proposition $\{c\}$ is supported by the initial conditions. Therefore, the consensus state reached by the CUE-based iteration is a refinement of $\widehat{GT}$. This is an extremely useful feature when sensor credibility is unavailable—the converged state does not contradict or invalidate the integrity of the existing knowledge base. A rational consensus can also be used to estimate sensor credibility [30].

2) Convergence With an Infinite Pool of Fusion Operators: In the absence of sensor reliability and/or credibility of evidence, an intuitive technique for belief revision involves weighting the incoming evidence based on its "distance" to the BoE that is updated, i.e., in updating $\mathcal{E}_i$, (in Definition 13) $\beta_{ij}(\cdot)$ are chosen s.t. $\lambda_j \propto \|\mathcal{E}_i - \mathcal{E}_j\|$. This would essentially lead to an infinite pool of fusion operators. However, as the iteration progresses, it is easy to see that such a pool approximates a finite pool where the $\lambda_{(\cdot)}$ of fusion operators corresponding to each agent are constant. Fig. 6 shows the convergence behavior of this scheme under two different $\widehat{GT}$s. Notice that in each case, the convergence time has increased, but the consensus state is still consistent with the $\widehat{GT}$.

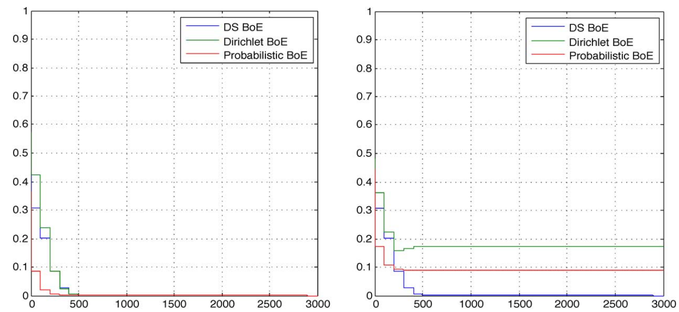

(a) w.r.t. consensus of approx. BoE (b) w.r.t. consensus of DS BoE

Fig. 7. **Comparison of convergence rates (as seen from $\|\mathcal{E}_1^{\cdot}[k] - \mathcal{E}^*\|$) for the DST BoE, Dirichlet BoE, and the pignistic p.m.f. models**. (**a**) Convergence rate is compared with the consensus of each model. (**b**) Converged rate is compared with the consensus the DST BoE. (a) w.r.t. consensus of approx. BoE; (b) w.r.t. consensus of DS BoE.

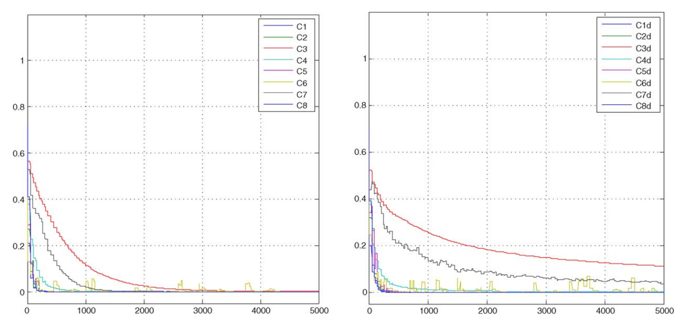

(a) **C1-C8** (b) **C1d-C8d**

Fig. 8. **Consensus formation (as seen from $\|\mathcal{E}_1[k] - \mathcal{E}^*\|$)**. (**a**) Synchronous configurations. (**b**) Asynchronous configurations. The small-world configurations (i.e., **C3**, **C7**, **C3d**, **C7d**) appear to take the longest time for convergence. (a) **C1–C8**; (b) **C1d–C8d**.

3) Convergence Rate Comparison: Here we compare the convergence rates under three complexity levels of uncertainty modeling: the general DST BoE versus Dirichlet BoE and the pignistic p.m.f. Given the agent BoE $\mathcal{E}_i$, we generate (**a**) the Dirichlet BoE $\hat{\mathcal{E}}_i$ by keeping all singleton masses the same and allocating all non-singleton masses to $\hat{m}_i(\Theta)$; and (**b**) the pignistic p.m.f. $\acute{\mathcal{E}}_i$ via the pignistic transformation (see Section III-C-5). In Fig. 7, the comparison is done with the consensus reached by each model. So, the final error necessarily approaches zero, although the consensus state of each model may not be identical.

C. Convergence in Complex Fusion Environments

Here, we use configurations **C1**–**C8** and **C1d**–**C8d** to study the convergence behavior of the proposed CUE-based belief revision under different agent interaction topologies.

Fig. 8 shows the variation of $\|\mathcal{E}_1[k] - \mathcal{E}^*\|$ for each configuration as the system evolves. Notice that the networks show somewhat similar pattern in convergence irrespective of the presence of communication delays. For both the synchronous and asynchronous cases, the small-world configurations (i.e., **C3**, **C7**, **C3d**, **C7d**)—which are characterized by high clustering and short path lengths—appear to take the longest time for convergence.

VII. Concluding Remarks

We examine the important problem of consensus formation among agents whose states are modeled as DST BoEs. Both asynchronous networks and situations where the link structure are dynamic are accounted for by utilizing the framework of asynchronous iterations of paracontracting operators. The DST agent model provides significant flexibility: p.m.f. agent models can be accommodated as a special case, and the DST model allows for capturing the types of imperfections and nuances that one encounters in soft evidence. So, this work is ideally suited to study opinion and consensus dynamics in social networks. An important consequence of the proposed DST belief revision is its ability to generate a rational consensus, implying that the consensus reached can be used to estimate agent credibility (even when the ground truth is partially known or unknown). The simulation results for a representative sample of common network configurations demonstrate the validity of the approach and the results obtained.

Appendix

A. Proof of Claim 5

1) Part 1: First, we need to show that there are infinitely many $\mathcal{E}^* \in \mathcal{E}_\Theta$ s.t. $\mathcal{E}^* = \mathrm{F}^i_{\triangleleft}(\mathcal{E}^*, \ldots, \mathcal{E}^*)$. To proceed, construct a BoE $\mathcal{E}^* \equiv \{\Theta, \mathcal{F}^*, m^*(\cdot)\}$, with $\mathcal{F}^*$ s.t. $\forall B \in \mathcal{F}^*$, $\nexists C \in \mathcal{F}^*$, satisfying $B \subset C$ or $C \subset B$. Such a BoE exists because $\mathcal{F}^* = \{B, \Theta \setminus B\}$ for any arbitrary $\emptyset \neq B \subset \Theta$ satisfies this condition. Now, we have $\mathcal{E}^* \in \mathcal{E}_\Theta$, for any arbitrary $m^*(\cdot) : 2^\Theta \mapsto [0,1]$ s.t. $\sum_{A \in \mathcal{F}^*} m^*(A) = 1$. Furthermore, by construction, we have $m^*(B|C) = 0$, for all $B, C \in \mathcal{F}^*$ s.t. $B \neq C$ and $m^*(B|B) = 1$, where both facts immediately follow from the CCT (Theorem 2).

Now, let $\mathcal{E}' = \mathrm{F}^i_{\triangleleft}(\mathcal{E}^*, \ldots, \mathcal{E}^*)$, where $\mathcal{E}' \equiv \{\Theta, \mathcal{F}', m'(\cdot)\}$. Then, for any $B \in \mathcal{F}^*$, we have

$$\begin{aligned} m'(B) &= \lambda_i m^*(B) + \sum_{j \neq i} \sum_{A \in \mathcal{F}^*} \lambda_j m^*(A) m^*(B|A) \\ &= \lambda_i m^*(B) + \sum_{j \neq i} \lambda_j m^*(B) m^*(B|B) = m^*(B). \end{aligned}$$

Therefore, we get $\mathcal{F}' = \mathcal{F}^*$ and $m'(B) = m^*(B)$, $\forall B \subseteq \Theta$. Hence $\mathcal{E}^* \in \mathrm{fix}(\mathrm{F}_{\triangleleft})$. However, $\mathcal{E}^* \in \mathcal{E}_\Theta$ is arbitrary and there are infinitely many $\mathcal{E}^* \in \mathcal{E}_\Theta$ satisfying the above construction. Hence, $\mathrm{F}^i_{\triangleleft}$ has infinitely many fixed points.

Part 2: Let $\mathcal{E}^* \in \mathrm{fix}(\mathrm{F}^i_{\triangleleft})$ be arbitrary. We need to show that $\nexists C \in \mathcal{F}^*$ s.t. $B \subset C$ or $B \supset C$, for all $B \in \mathcal{F}^*$. Now, the mass function for any arbitrary $B \in \mathcal{F}_\Theta$ corresponding to $\mathcal{E}^* = \mathrm{F}^i_{\triangleleft}(\mathcal{E}^*, \ldots, \mathcal{E}^*)$ is given by

$$\begin{aligned} m^*(B) &= \lambda_i m^*(B) + \sum_{j \neq i} \sum_{A \in \mathcal{F}^*} \lambda_j m^*(A) m^*(B|A) \\ &= \lambda_i m^*(B) + (1 - \lambda_i) \sum_{A \in \mathcal{F}^*} m^*(A) m^*(B|A). \end{aligned}$$

This implies that $m^*(B) = \sum_{A \in \mathcal{F}^*} m^*(A) m^*(B|A)$, since $\lambda_i \neq 1$. For any $B \in \mathcal{F}^*$ s.t. $\nexists C \in \mathcal{F}^*$ satisfying $C \subset B$, we get

$$\begin{aligned} m^*(B) &= m^*(B) m^*(B|B) + \sum_{B \neq A \in \mathcal{F}^*} m^*(A) m^*(B|A) \\ \implies & \sum_{B \neq A \in \mathcal{F}^*} m^*(A) m^*(B|A) = 0, \text{ since } m^*(B|B) = 1. \end{aligned}$$

So, $m^*(B|A) = 0$, $\forall A \in \mathcal{F}^*$, s.t. $A \neq B$. The result follows from the fact that $A \not\supseteq B$ for all $A \in \mathcal{F}^*$ s.t. $A \neq B$ according to the CCT. This completes the proof of Claim 5. ■

B. Proof of Claim 6

Consider a set of BoEs $\mathcal{E}_j \in \mathcal{E}_\Theta, j = 1, \ldots, m$, representing the agent states after sufficiently large number of iterations s.t. $\mathcal{F}_i = \cup_j \mathcal{F}_j$ for $\mathcal{E}_i \neq \hat{\mathcal{E}}^t$ and $\mathcal{F}_i = \{B \in \mathcal{F}_j | \exists C \in \mathcal{F}_i \text{ s.t. } B \subseteq C, j = 1, \ldots, m\}$ for all others.

Let $\mathcal{E}^* \in \mathrm{fix}(\mathrm{F}^i_{\triangleleft})$ be arbitrary. We need to prove that $\|\mathrm{F}^i_{\triangleleft}(\mathcal{E}_1, \ldots, \mathcal{E}_m) - \mathcal{E}^*\| < \max_j \|\mathcal{E}_j - \mathcal{E}^*\|$, or otherwise $\mathcal{E}_j = \mathcal{E}^*, j = 1, \ldots, m$. Now, the claim is trivial for $\mathcal{E}_j = \mathcal{E}^*, j = 1, \ldots, m$. Let us proceed as follows to show the other case.

Let $\mathcal{B}_j = \{B \in \mathcal{F}_j | \nexists C \in \mathcal{F}_j \text{ s.t. } C \subset B\}$, for $j = 1, \ldots, m$. For any $B \in \mathcal{B}_j$,

$$\begin{aligned} &\sum_{A \in \mathcal{F}_j} m_j(A) m_j(B|A) \\ &\quad = \sum_{B \subseteq A \in \mathcal{F}_j} m_j(A) m_j(B|A) + \sum_{B \supset A \in \mathcal{F}_j} m_j(A) \underbrace{m_j(B|A)}_{=0} \\ &\quad = m_j(B) \underbrace{m_j(B|B)}_{=1} + \sum_{B \subset A \in \mathcal{F}_j} m_j(A) m_j(B|A) \\ &\implies \sum_{A \in \mathcal{F}_j} m_j(A) m_j(B|A) \geq m_j(B), \quad \forall B \in \mathcal{B}_j. \end{aligned} \tag{13}$$

Since $\sum_{B \subseteq \Theta} \sum_{A \in \mathcal{F}_j} m_j(A) m_j(B|A) = 1$, we must also have

$$\sum_{A \in \mathcal{F}_j} m_j(A) m_j(B|A) \leq m_j(B), \quad \forall B \notin \mathcal{B}_j. \tag{14}$$

Now, consider $\mathcal{F}^*$. From Claim 5, we have $\nexists C \in \mathcal{F}^*$ s.t. $B \subset C$ or $B \supset C$, for all $B \in \mathcal{F}^*$. Therefore, from (13), we must have

$$\sum_{B \in \mathcal{F}^*} \left| \sum_{A \in \mathcal{F}_j} m_j(A) m_j(B|A) - m^*(B) \right|^p < \sum_{B \in \mathcal{F}^*} |m_j(B) - m^*(B)|^p, \tag{15}$$

for any $\mathcal{E}_j \neq \mathcal{E}^*$ and $p \in \Re_+$ s.t. $p \geq 1$, since $\sum_{B \in \mathcal{F}^*} m^*(B) = 1$. Also, for all $B \notin \mathcal{F}^*$,

$$\begin{aligned} &\sum_{B \notin \mathcal{F}^*} \left| \sum_{A \in \mathcal{F}_j} m_j(A)\, m_j(B|A) - m^*(B) \right|^p \\ &\quad = \sum_{B \notin \mathcal{F}^*} \left| \sum_{A \in \mathcal{F}_j} m_j(A) m_j(B|A) \right|^p, \quad \text{since } m^*(B) = 0 \\ &\quad \leq \sum_{B \notin \mathcal{F}^*} |m_j(B)|^p, \quad \text{from (14)} \\ &\quad = \sum_{B \notin \mathcal{F}^*} |m_j(B) - m^*(B)|^p, \quad \text{since } m^*(B) = 0. \end{aligned}$$

Therefore,

$$\sum_{B \notin \mathcal{F}^*} \left| \sum_{A \in \mathcal{F}_j} m_j(A) m_j(B|A) - m^*(B) \right|^p \leq \sum_{B \notin \mathcal{F}^*} |m_j(B) - m^*(B)|^p. \tag{16}$$

From (15) and (16), we get

$$\sum_{B\subseteq\Theta}\left|\sum_{A\in\mathcal{F}_j} m_j(A)m_j(B|A)-m^*(B)\right|^p < \sum_{B\subseteq\Theta}|m_j(B)-m^*(B)|^p = \|\mathcal{E}_j-\mathcal{E}^*\|^p. \quad (17)$$

Similarly, we can also show that

$$\sum_{B\subseteq\Theta}\left|\sum_{A\in\mathcal{F}_j} m_i(A)m_j(B|A)-m^*(B)\right|^p < \|\mathcal{E}_i-\mathcal{E}^*\|^p. \quad (18)$$

Now, for an update of $\mathcal{E}_i \neq \hat{\mathcal{E}}^t$, use (17) to show that

$$\begin{aligned}
&\left\|\mathrm{F}^i_{\triangleleft}(\mathcal{E}_1,\ldots,\mathcal{E}_m)-\mathcal{E}^*\right\|^p \\
&=\sum_{B\subseteq\Theta}\left|\lambda_i m_i(B)+\sum_{j\neq i}\sum_{A\in\mathcal{F}_j}\lambda_j m_j(A)m_j(B|A)-m^*(B)\right|^p \\
&\leq \lambda_i\sum_{B\subseteq\Theta}|m_i(B)-m^*(B)|^p \\
&\quad+\sum_{j\neq i}\lambda_j\sum_{B\subseteq\Theta}\left|\sum_{A\in\mathcal{F}_j}m_j(A)m_j(B|A)-m^*(B)\right|^p \\
&<\lambda_i\|\mathcal{E}_i-\mathcal{E}^*\|^p+\sum_{j\neq i}\lambda_j\|\mathcal{E}_j-\mathcal{E}^*\|^p, \quad \text{from (17)} \\
&=\sum_j\lambda_j\|\mathcal{E}_j-\mathcal{E}^*\|^p \\
&\leq\sum_j\lambda_j\max_k\|\mathcal{E}_k-\mathcal{E}^*\|^p=\max_k\|\mathcal{E}_k-\mathcal{E}^*\|^p.
\end{aligned}$$

This implies that $\|\mathrm{F}^i_{\triangleleft}(\mathcal{E}_1,\ldots,\mathcal{E}_m)-\mathcal{E}^*\| < \max_k\|\mathcal{E}_k-\mathcal{E}^*\|$.

Similarly, use (18) to show that $\|\mathrm{F}^i_{\triangleleft}(\mathcal{E}_1,\ldots,\mathcal{E}_m)-\mathcal{E}^*\| < \max_k\|\mathcal{E}_k-\mathcal{E}^*$, holds true for an update of $\mathcal{E}_i=\hat{\mathcal{E}}^t$. This completes the proof.

Acknowledgment

The discussions with and assistance of Messrs. Ranga Dabarera and Jian Xu are gratefully appreciated.

References

[1] K. Lehrer and C. Wagner, *Rational Consensus in Science and Society*, ser. Philosophical studies series in philosophy. Dordrecht, The Netherlands: D. Reidel, 1981.

[2] H. E. Stephanou and S.-Y. Lu, "Measuring consensus effectiveness by a generalized entropy criterion," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 10, no. 4, pp. 544–554, Jul. 1988.

[3] M. W. Macy and R. Willer, "From factors to actors: Computational sociology and agent-based modeling," *Annu. Rev. Sociol.*, vol. 28, pp. 143–166, Aug. 2002.

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

[5] W. Ren, R. W. Beard, and D. B. Kingston, "Multi-agent Kalman consensus with relative uncertainty," in *Proc. Amer. Control Conf. (ACC)*, Portland, OR, Jun. 2005, pp. 1865–1870.

[6] L. Xiao, S. Boyd, and S. Lall, "A scheme for robust distributed sensor fusion based on average consensus," in *Proc. Int. Symp. Inf. Process. Sens. Netw. (IPSN)*, Apr. 2005, pp. 63–70.

[7] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.

[8] W. Ren, R. W. Beard, and E. M. Atkins, "Information consensus in multivehicle cooperative control," *IEEE Control Syst. Mag.*, vol. 27, no. 2, pp. 71–82, Apr. 2007.

[9] L. Schenato and G. Gamba, "A distributed consensus protocol for clock synchronization in wireless sensor network," in *Proc. IEEE Conf. Decision Control (CDC)*, Dec. 2007, pp. 2289–2294.

[10] M. Cao, A. S. Morse, and B. D. O. Anderson, "Reaching a consensus in a dynamically changing environment: A graphical approach," *SIAM J. Control Optimiz.*, vol. 47, no. 2, pp. 575–600, 2008.

[11] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise," *IEEE Trans. Signal Process.*, vol. 57, no. 1, pp. 355–369, Jan. 2009.

[12] Y. G. Sun and L. Wang, "Consensus of multi-agent systems in directed networks with nonuniform time-varying delays," *IEEE Trans. Autom. Control*, vol. 54, no. 7, pp. 1607–1613, Jul. 2009.

[13] R. F. Weiss, "Consensus technique for the variation of source credibility," *Psychol. Rep.*, vol. 20, no. 3c, pp. 1159–1162, Jan. 2011.

[14] Z. Wu, H. Fang, and Y. She, "Weighted average prediction for improving consensus performance of second-order delayed multi-agent systems," *IEEE Trans. Syst., Man, Cybern., B: Cybern.*, vol. 42, no. 5, pp. 1501–1508, Oct. 2012.

[15] F. Borran, M. Hutle, N. Santos, and A. Schiper, "Quantitative analysis of consensus algorithms," *IEEE Trans. Depend. Secure Comput.*, vol. 9, no. 2, pp. 236–249, Mar. 2012.

[16] C. Altafani, "Consensus problems on networks with antagonistic interactions," *IEEE Trans. Autom. Control*, vol. 58, no. 4, pp. 935–946, Apr. 2013.

[17] Z. Li, W. Ren, X. Liu, and M. Fu, "Consensus of multi-agent systems with general linear and Lipschitz nonlinear dynamics," *IEEE Trans. Autom. Control*, vol. 58, no. 7, pp. 1786–1791, Jul. 2013.

[18] X. Lu, R. Lu, S. Chen, and J. Lu, "Finite-time distributed tracking control for multi-agent systems with a virtual leader," *IEEE Trans. Circuits Syst. I*, vol. 60, no. 2, pp. 352–362, Feb. 2013.

[19] H. Terelius, D. Varagnolo, C. Baquero, and K. H. Johansson, "Fast distributed estimation of empirical mass functions over anonymous networks," in *Proc. IEEE Conf. Decision Control (CDC)*, Florence, Italy, Dec. 2013, pp. 6771–6777.

[20] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava, "Participatory sensing," in *Proc. Workshop World-Sensor-Web (WSW): Mobile Device Centric Sens. Netw. Applicat.*, 2006, pp. 117–134.

[21] G. Shafer, *A Mathematical Theory of Evidence*. Princeton, NJ, USA: Princeton Univ. Press, 1976.

[22] 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. Int. Conf. Inf. Fusion (FUSION)*, Seattle, WA, USA, Jul. 2009, pp. 2122–2129.

[23] 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 Theory of Functions (BELIEF)*, Brest, France, Apr. 2010.

[24] T. L. Wickramarathne, K. Premaratne, M. N. Murthi, M. Scheutz, and S. Kübler, "Belief theoretic methods for soft and hard data fusion," in *Proc. Int. Conf. Statist. Signal Process. (ICASSP)*, Prague, Czech Republic, May 2011.

[25] T. L. Wickramarathne, K. Premaratne, and M. N. Murthi, "Toward efficient computation of the dempster-shafer belief theoretic conditionals," *IEEE Trans. Cybern.*, vol. 43, no. 2, pp. 712–724, Apr. 2012.

[26] T. L. Wickramarathne, K. Premaratne, and M. N. Murthi, "Convergence analysis of consensus belief functions within asynchronous ad-hoc fusion networks," in *Proc. Int. Conf. Statist. Signal Process. (ICASSP)*, Vancouver, BC, Canada, May 2013.

[27] M. Pott, "On the convergence of asynchronous iteration methods for nonlinear paracontractions and consistent linear systems," *Linear Algebra and Its Applicat.*, vol. 283, no. 13, pp. 1–33, 1998.

[28] R. Fagin and J. Y. Halpern, "Uncertainty, belief and probability," *Comput. Intell.*, vol. 7, pp. 160–173, 1991.

[29] R. Fagin and J. Y. Halpern, "A new approach to updating beliefs," in *Proc. Conf. Uncertainty Artif. Intell. (UAI)*, P. P. Bonissone, M. Henrion, L. N. Kanal, and J. F. Lemmer, Eds., 1991, pp. 347–374, New York, NY: Elsevier Science.

[30] 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.

[31] L. Fang and P. J. Antsaklis, "Asynchronous consensus protocols using nonlinear paracontractions theory," *IEEE Trans. Autom. Control*, vol. 53, no. 10, pp. 2351–2355, Nov. 2008.
[32] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Trans. Autom. Control*, vol. 50, no. 2, pp. 169–182, Feb. 2005.
[33] D. B. West, *Introduction to Graph Theory*, 2nd ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2001.
[34] T. L. Wickramarathne, K. Premaratne, and M. N. Murthi, "Focal elements generated by the Dempster-Shafer theoretic conditionals: A complete characterization," in *Proc. Int. Conf. Inf. Fusion (FUSION)*, Edinburgh, U.K., Jul. 2010.
[35] P. Vannoorenberghe, "On aggregating belief decision trees," *Inf. Fusion*, vol. 5, no. 3, pp. 179–188, Sep. 2004.
[36] K. Premaratne, D. A. Dewasurendra, and P. H. Bauer, "Evidence combination in an environment with heterogeneous sources," *IEEE Trans. Syst., Man, Cybern., A: Syst. Humans*, vol. 37, no. 3, pp. 298–309, May 2007.
[37] P. Smets, "Practical uses of belief functions," in *Proc. Conf. Uncertainty Artif. Intell. (UAI)*, K. B. Laskey and H. Prade, Eds., 1999, pp. 612–621, San Francisco, CA: Morgan Kaufmann.
[38] G. Kossinets, J. Kleinberg, and D. Watts, "The structure of information pathways in a social communication network," in *Proc. ACM SIGKDD Int. Conf. Knowl. Disc. Data Mining (KDD)*, New York, NY, USA, 2008, pp. 435–443.
[39] P. Erdös and A. Rényi, "On the evolution of random graphs," *Inst. Math., Hungarian Acad. Sci.*, vol. 5, pp. 17–61, 1960.
[40] D. J. Watts and S. H. Strogatz, "Collective dynamics of "small-world" networks," *Nature*, vol. 393, no. 6684, pp. 440–442, Jun. 1998.
[41] A.-L. Barabási, R. Albert, and H. Jeong, "Scale-free characteristics of random networks: The topology of the world-wide web," *Physica A: Statist. Mech. its Applicat.*, vol. 281, no. 1–4, pp. 69–77, Jun. 2000.
[42] E. Ravasz and A.-L. Barabási, "Hierarchical organization in complex networks," *Phys. Rev. E*, vol. 67, no. 2, p. 026112, Feb. 2003.
[43] M. Webster, *Merriam Websters Collegiate Dictionary*. Springfield, MA, USA: Merriam Webster, Jul. 2003.
[44] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Rev. Modern Phys.*, vol. 74, no. 1, pp. 47–97, Jan. 2002.

Thanuka L. Wickramarathne (M'12) received his B.Sc. degree in Electronics and Telecommunication Engineering from University of Moratuwa, Sri Lanka, in 2006 and his M.S. and Ph.D. degrees in Electrical and Computer Engineering, both from University of Miami, Coral Gables, FL, in 2008 and 2012 respectively.

Presently he is a Research Assistant Professor in the Department of Computer Science and Engineering at the University of Notre Dame, Notre Dame, Indiana, USA. His research interests are in the general areas of multi-sensor data fusion, statistical signal processing, knowledge discovery in uncertain data domains, belief theory, sensor networks, complex networks and probabilistic graphical models.

Kamal Premaratne (SM'94) received the B.Sc. degree in Electronics and Telecommunication Engineering (with First-Class Honors) in 1982 from University of Moratuwa, Sri Lanka. He obtained his M.S. and Ph.D. degrees, both in Electrical and Computer Engineering, in 1984 and 1988 respectively, from the University of Miami, Coral Gables, Florida, USA, where he is presently a Professor.

He has received the 1992/93 "Mather Premium" and 1999/00 "Heaviside Premium" of the Institution of Electrical Engineers (IEE), London, UK, and the 1991, 1994 and 2001 "Eliahu I. Jury Excellence in Research Award" of the College of Engineering, University of Miami. He has served as an Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING (1994–1996) and the *Journal of the Franklin Institute* (1993–2005). He is a Fellow of IET (formerly IEE). His current research interests include belief theory, evidence fusion, machine learning and knowledge discovery from imperfect data, and network congestion control.

Manohar N. Murthi (M'94) received his B.S. degree in electrical engineering and computer science from the University of California, Berkeley, in 1990, and his M.S. and Ph.D. degrees in electrical engineering (communication theory and systems) from the University of California, San Diego, CA, in 1992 and 1999, respectively.

He has previously worked at Qualcomm in San Diego, CA, KTH (Royal Institute of Technology), Stockholm, Sweden, and Global IP Sound in San Francisco, CA. In September 2002 he joined the Department of Electrical and Computer Engineering, University of Miami, Coral Gables, FL, where he is an Associate Professor. His research interests are in the general areas of signal and data modeling, compression, fusion and learning, and networking. He is a recipient of a National Science Foundation CAREER Award.

Nitesh V. Chawla (M'02) received his B.E. degree in Computer Science, University of Pune, India, in 1997 and his M.S. and Ph.D. degrees in Computer Science and Engineering, both from University of South Florida, FL, in 1999 and 2002 respectively. He is currently the Frank Freimann Collegiate Associate Professor of Computer Science and Engineering and the Director of Interdisciplinary Center for Network Science and Applications (iCeNSA) at the University of Notre Dame, Notre Dame, IN, USA. His research interests are in data and network science, and various interdisciplinary applications such as healthcare, social networks, and environmental sciences.