Chapter 15

A NETWORKED EVIDENCE THEORY FRAMEWORK FOR CRITICAL INFRASTRUCTURE MODELING

Chiara Foglietta, Andrea Gasparri and Stefano Panzieri

Abstract   This paper describes a distributed approach for data fusion and information sharing based on evidence theory and the transferable belief model. Evidence theory aggregates data generated from different sources in order to better assess an ongoing situation and to aid in the response and decision making processes. In the domain of critical infrastructure protection, researchers are forced to develop distributed approaches for modeling and control with a minimal exchange of data due to the existence of multiple stakeholders and interconnections between infrastructure components. Evidence theory permits the modeling of uncertainty in data fusion, but it is typically applied in a centralized manner. This paper proposes a decentralized extension of the transferable belief model that facilitates the application of evidence theory to data fusion in critical infrastructure applications. A case study is provided to demonstrate the convergence of results similar to the centralized approach, and to show the utility of fusing data in a distributed manner for interdependent critical infrastructure systems.

Keywords: Modeling, evidence theory, situational awareness, data fusion

1.   Introduction

A nation's critical infrastructure comprises complex, interdependent systems whose proper operation and interaction are essential to the welfare of society. Modeling the interdependencies between the different critical infrastructure sectors is a complex task, but this is vital to correctly analyze and predict cascading phenomena.

Interdependencies, from the viewpoint of data fusion, can be analyzed to discern critical events and failures. Events associated with critical infrastructures typically have low probabilities but high impact. The impact is exacerbated

by effects that are not localized and tend to propagate to interconnected infrastructures. This implies that a "signature" of the event can be identified in the interconnected system and, sometimes, in independent agencies such as meteorological services or police departments. Therefore, there is a need to fuse together the available data and derive a common belief of the current situation. Developing a common belief facilitates efficient and effective decisions and plans of action. Since critical infrastructures are inherently distributed [5], the fusion mechanism should also be distributed.

This paper provides an extension of the transferable belief model to facilitate the application of evidence theory. The existing transferable belief model provides a centralized technique for fusing data. However, a decentralized approach is required for interdependent critical infrastructures. This paper demonstrates the application of an extended decentralized transferable belief model to the domain of critical infrastructure protection and shows how data fusion can help develop more accurate beliefs about interdependent infrastructures.

2.   Background

Evidence theory is a methodology that is commonly applied to data fusion problems. Data fusion seeks to combine data from heterogeneous sources or sensors to provide estimates of ongoing events [3].

Evidence theory stems primarily from the pioneering work of Dempster [4] and Shafer [10] and is often considered in the same light as Bayesian networks. The primary difference, however, is the ability to deal with uncertainty [9]. The Dempster-Shafer theory [10] explicitly considers uncertainty and examines if the various sources of data are inconsistent or if there is an error in the modeling process. On the other hand, Bayesian networks use the recognition of input values as a likelihood from pre-determined patterns. The application of the two approaches depends on the type of knowledge that has to be represented and fused [8].

Several methods have been proposed for combining and correlating the available data in the context of the Dempster-Shafer framework. This work uses the methodology proposed by Smets [11], which extends the Dempster-Shafer framework by assuming that the correct answer might not be among the considered ones (i.e., it engages the open world assumption). Smets' approach also allows the computation of the amount of contradictory data in the value of the empty set.

The major limitation to applying evidence theory in a real context is the number of hypotheses required to model the application of interest. This can be explained by the fact that, from a computational perspective, the power set of the set of hypotheses has to be computed, causing the complexity to grow exponentially with the number of hypotheses. However, evidence theory has been successfully applied to a variety of practical problems using approximations (see, e.g., [7, 13]).

Data fusion can also aid in impact assessment. In fact, limiting an impact assessment strictly to measured events can lead to heavy underestimation or in-

correct estimation. For example, an isolated failure can propagate in a number of ways depending on the cause and the detection methods. Examples include a fire blast detected by a sensor and a computer virus detected by an intrusion detection system. In the first example, the fire blast propagation effects are associated with an interdependency model according to a spatial proximity pattern. In the second example, a computer virus may propagate to similar, and highly dispersed, telecommunication nodes.

3.   Evidence Theory

Evidence theory provides a means to form a consolidated belief by correlating evidence from different sources [4, 10, 11]. This section provides an overview of the evidence theory formalisms used in this work.

Let ω_i represent a cause of system failure and $\Omega = \{\omega_1, \cdots, \omega_n\}$ be the set of hypotheses containing known possible failures. This set is called the "frame of discernment." For example, the possible causes of failures of a critical infrastructure asset include sabotage, device failure, fault due to weather and a (cyber) denial-of-service attack. Note that the hypotheses are assumed to be mutually exclusive in evidence theory.

The power set of the frame of discernment is expressed as $\Gamma(\Omega) = \{\gamma_1, \cdots, \gamma_{2^{|\Omega|}}\}$, which has cardinality $|\Gamma(\Omega)| = 2^{|\Omega|}$. The power set contains all possible subsets of Ω, including the empty set $\gamma_1 = \emptyset$ and the universal set $\gamma_{2^{|\Omega|}} = \Omega$.

The transferable belief model [11] is derived from the basic belief mass function m:

$$m : \Gamma(\Omega) \rightarrow [0, 1].$$

This function, also called the "basic belief assignment" (BBA), maps each element of the power set to a value between 0 and 1. Each BBA is an atomic element in the transferable belief model. In fact, each sensor, agent or node must be able to assign the BBA values by some subjective assumptions or through algorithms that automatically determine the assignment. The BBA function is constrained by:

$$\sum_{\gamma_a \subseteq \Gamma(\Omega)} m(\gamma_a) = 1 \quad \text{with} \quad m(\emptyset) = 0.$$

The transferable belief model examines propositions accordingly as: "the true value of ω_i is in γ_a" where $\gamma_a \in \Gamma(\Omega)$. For $\gamma_a \in \Gamma(\Omega)$, $m(\gamma_a)$ is the confidence that supports exactly γ_a. This implies that the true value is in the set γ_a; however, due to the lack of additional data, it is not possible to support any strict subset of γ_a.

In the case of different independent data sources, a rule is necessary to aggregate the data. Several rules of combination exist in the literature; the most widely used rules are Dempster's rule [4] and Smets' rule [11].

Dempster's rule of combination [4], which was the first to be proposed, is a purely conjunctive operation. The rule strongly emphasizes the agreement between multiple sources and ignores conflicting evidence using a normalization

factor as shown in the following equation:

$$\text{Dempster}\{m_i, m_j\}(\emptyset) = 0$$

$$\text{Dempster}\{m_i, m_j\}(\gamma_a) = \frac{\sum\limits_{\gamma_b \cap \gamma_c = \gamma_a} m_i(\gamma_b) m_j(\gamma_c)}{1 - \sum\limits_{\gamma_b \cap \gamma_c = \emptyset} m_i(\gamma_b) m_j(\gamma_c)} \qquad \forall \gamma_a \in \Gamma(\Omega).$$

On the other hand, Smets' rule of combination [11] provides the ability to explicitly express the contradiction in the transferable belief model by letting $m(\emptyset) \neq 0$. This combination rule, unlike Dempster's rule, avoids normalization while preserving commutativity and associativity:

$$\text{Smets}\{m_i, m_j\}(\gamma_a) = m_i(\gamma_a) \otimes m_j(\gamma_a) \qquad \forall \gamma_a \in \Gamma(\Omega)$$

where

$$m_i(\gamma_a) \otimes m_j(\gamma_a) = \sum_{\gamma_b \cap \gamma_c = \gamma_a} m_i(\gamma_b) m_j(\gamma_c) \qquad \forall \gamma_a \in \Gamma(\Omega).$$

The relation $m(\emptyset) > 0$ can be explained in two ways: (i) open world assumption; and (ii) quantified conflict. The open world assumption, proposed by Dempster, reflects the idea that the frame of discernment must contain the true value. If the open world assumption is true, then the set of hypotheses must contain all possibilities. Under this interpretation, if $\emptyset$ is the complement of Ω, the mass $m(\emptyset) > 0$ represents the case where the truth is not contained in Ω. Alternatively, the notion of quantified conflict means that there is some underlying conflict between the sources that are combined to produce the BBA. Hence, the mass assigned to $m(\emptyset)$ represents the degree of conflict. Specifically, it is computed as:

$$m_i(\emptyset) \otimes m_j(\emptyset) = 1 \;\; - \sum_{\gamma_a \in \Gamma, \gamma_a \neq \emptyset} (m_i(\gamma_a) \otimes m_j(\gamma_a)).$$

4.   Data Fusion in the Network Context

Consider a network of multiple agents described by an indirect graph $\mathcal{G} = \{V, E\}$ where $V = \{v_i \mid i = 1, \cdots, n_V\}$ is the set of nodes and $E = \{e_{ij} \mid (v_i, v_j)\}$ is the set of edges that represent a communication channel between the nodes. Edges are indirect and, thus, the existence of an arc e_{ij} implies the existence of an edge e_{ji}.

In this work, we assume that no central unit is available to perform data aggregation. Additionally, communications between nodes are limited to the neighbors of the node under consideration (i.e., nodes that are physically or directly connected to the node under consideration). These assumptions are reasonable for data fusion problems in the area of sensor networks.

Table 1. BBA assignments for a telecommunications network.

Set	Node 1	Node 2	Node 3	m_{12}	m_{123}
$\emptyset$	0.0	0.0	0.0	0.44	0.770
{a}	0.1	0.5	0.7	0.11	0.095
{b}	0.8	0.4	0.2	0.44	0.134
{a,b}	0.1	0.1	0.1	0.01	0.010

A direct consequence of the assumptions, however, is that Smets' rule of composition cannot be directly applied. Indeed, applying Smets' rule multiple times over the same BBAs leads to different outputs. Note that this could easily happen in a distributed context where communications are local and limited to the one-hop neighborhood.

Consider, for example, a scenario where it is necessary to identify the degradation of services in a telecommunications network. In the case of extensive delays in packet transmission, it may be necessary to determine if the situation is a temporary congestion or a persistent malicious attack. Network sensors (e.g., intrusion detection systems) can detect a denial-of-service (DoS) attack. The DoS attack may result in cascading effects within the network that are identified by other sensors. If the attack is conducted on a sufficiently large scale, entire regions can be compromised, with different sensors providing different inputs based on localized knowledge.

A general method to define a BBA allocation is to consider the reliability of the data source. Suppose that the data obtained from the source supports a set of hypotheses in Ω. Then, the subset γ_a of the power set $\Gamma(\Omega)$, containing the set of hypotheses, receives a mass $m(\gamma_a)$ equal to the reliability of the source. The remaining mass $1 - m(\gamma_a)$ is assigned to the universal set because no other data is available.

Table 1 shows the BBA assignments for a cause classification problem in a telecommunications network using the centralized approach. The network has three sources (nodes) that relay data to determine the probable causes of a network congestion incident. The table shows the BBA values assigned to the various network nodes. The frame of discernment is $\Omega = \{a, b\}$, where hypothesis a is a denial-of-service attack (i.e., congestion of one or more network nodes that intentionally degrades the telecommunications network), and hypothesis b indicates congestion in the telecommunications network due to routing problems.

If we assume that Node 1 is coordinating with Node 2, then upon applying Smets' operator, the result is:

$$m_1 \otimes m_2 = \{0.44, 0.11, 0.44, 0.01\}.$$

Algorithm 1 : Gossip Algorithm

Data: $t = 0$, $s_i(t = 0)$ $\forall i = 1, \cdots, n$
Results: $s_i(t_{end})$ $\forall i = 1, \cdots, n$

while *end_condition* **do**
Select an edge $e_{ij} \in E(t)$ according to **e**;
Update the states of the selected nodes by applying the operator $\mathcal{R}$:
$s_i(t+1) = s_i(t) \otimes s_j(t)$
$s_j(t+1) = s_j(t) \otimes s_i(t)$
Set $t = t + 1$
end

If Node 1 then communicates with Node 3 by exchanging data about the possible cause of the congestion, then the result obtained by centralized aggregation is equal to:

$$m_{12} \otimes m_3 = \{0.77, 0.095, 0.134, 0.001\}.$$

Next, we consider the communications between Node 1 and Node 3 by applying Smets' operator. The result, which is different from the one obtained in the centralized system, is given by:

$$m_{123} \otimes m_3 = \{0.8828, 0.0767, 0.0404, 0.0001\}.$$

If a decision on the cause of the fault has to be made, then the latter case results in a change of opinion; according to the latest aggregations, the cause is a denial-of-service attack (hypothesis a) instead of network routing congestion (hypothesis b).

As shown, Smets' rule of combination cannot be directly used in a distributed data aggregation context. The next section presents a distributed algorithm that can update the knowledge of all the nodes in a network.

5.   Data Fusion Algorithm

The algorithm proposed by Gasparri, *et al.* [6] provides the ability to divide the knowledge of each node into two parts: (i) data shared between two nodes; and (ii) localized data retained by each node.

In the decentralized algorithm, the network is described by an indirect graph $\mathcal{G} = \{V, E\}$. A spanning tree $\mathcal{T} = \{V, \widehat{E}\}$ is derived, where $\widehat{E} \subseteq E$ is available to all nodes. Note that the nodes must have the capacity to save data. Interested readers are referred to [2] for details about spanning tree construction.

Communications between nodes are asynchronous and follow the Gossip Protocol [1]. This protocol is formalized as Algorithm 1. The triplet $\{\mathcal{S}, \mathcal{R}, \mathbf{e}\}$ specifies the network such that:

- $\mathcal{S}$ is the set of local states of each node in the network.

- $\mathcal{R}$ is local interaction rule, i.e., for each pair of nodes (i, j) such that $e_{ij} \in E$, the following equation holds:

$$\mathcal{R} : \mathbb{R}^q \times \mathbb{R}^q \to \mathbb{R}^q$$

where q is the number of elements called "focal sets."

- $\mathbf{e}$ is the process of edge selection for which $e_{ij} \in E(t)$ is the selected edge at time t.

To define the operator $\mathcal{R}$, it is first necessary to introduce the operator $\odot$. Consider two sets of BBAs:

$$m_k = \{m_k(\gamma_a) \mid \forall \gamma_a \in \Gamma(\Omega)\}$$

and

$$m_i = \{m_i(\gamma_a) \mid \forall \gamma_a \in \Gamma(\Omega)\},$$

such that $m_k = m_i \otimes m_j$. The operator $\odot$ can then be defined as:

$$m_j = m_k \odot m_i \triangleq \tilde{m}_k^i.$$

Starting with the power set element with the highest cardinality and examining elements with decreasing cardinality, the value of a BBA can be computed recursively as follows:

$$m_j(\gamma_a) = \frac{m_k(\gamma_a) \; - \displaystyle\sum_{\gamma_b \cap \gamma_c = \gamma_a, \gamma_a \subset \gamma_b} m_j(\gamma_b) m_i(\gamma_c)}{\displaystyle\sum_{\gamma_a \subseteq \gamma_b} m_i(\gamma_b)}.$$

It is now possible to introduce the operator $\mathcal{R}$, along with $\oplus$, to aggregate BBAs of the nodes:

$$m_i(t+1) = m_j(t+1) = m_i(t) \oplus m_j(t)$$
$$= \{\left(\tilde{m}_i^j(t, \gamma_a) \otimes \tilde{m}_i^j(t, \gamma_a)\right) \otimes \bar{m}_{i,j}(t, \gamma_a), \forall \gamma_a \in \Gamma(\Omega)\}.$$

Note that the term $\tilde{m}_i^j(t, \gamma_a)$ indicates the innovation of node i with respect to the node j, which can be calculated recursively using the operator $\mathcal{S}$:

$$\tilde{m}_i^j(t, \gamma_a) = m_i(t, \gamma_a) \odot \bar{m}_{i,j}(t, \gamma_a).$$

The element $\bar{m}_{i,j}(t, \gamma_a)$ express the common knowledge (i.e., knowledge exchanged between two agents (i, j) after the last aggregation) such that:

$$\bar{m}_{i,j}(t, \gamma_a) = \mathbf{n} = \{0, 0, \cdots, 0, 1\}.$$

Gasparri, *et al.* [6] have shown that the algorithm converges to the same result as the centralized aggregation algorithm. The convergence time of the distributed algorithm is related to the diameter d of the spanning tree. Additionally, the computational complexity of the $\mathcal{R}$-operator is the same as that of Smets' operator.

6.   Application Scenario

Applying the transferable belief model involves modeling a problem, specifying the frame of discernment and selecting the BBAs. Note that the assignment of BBAs is problem dependent and depends significantly on the source of information and knowledge about the system. In this work, we assume that experts provide advice on assigning the BBAs. Other possible approaches are described in [3, 12], where the reliability of the data sources is considered along with the effect of mass assignment on the compound hypotheses.

Our scenario involves a supervisory control and data acquisition (SCADA) system that aggregates data from a large number sensors positioned at remote locations. Operators located at a control center manage operations by acquiring system data and updating system parameters. Note that the remote facilities all have data regarding critical events and that the data is generated from different sources.

Communications between the control center and the devices in the field occur via dedicated telecommunications circuits or shared public media such as the Internet. Note that the same communications channels can support information sharing among critical infrastructures to help discern the possible causes of failures. The concept of information sharing across sectors can help prevent cascading effects or prevent the impact from propagating to interconnected infrastructures that have not yet been affected.

Our scenario considers $n = 5$ interdependent critical infrastructures. Each infrastructure owner determines the BBAs in his/her infrastructure. To discern the cause of a fault, the BBAs have to be aggregated and shared among the five infrastructures. The frame of discernment is $\Omega = \{a, b, c\}$, where hypothesis a represents a cyber attack, hypothesis b represents the failure of an isolated single unit, and hypothesis c represents a natural disaster (e.g., earthquake). Table 2 shows the BBA assignments for the five infrastructure nodes.

First, we evaluate the scenario using the centralized algorithm. The results as shown in Table 3. Each column presents the results obtained using Smets' operator, with column "NN 12" representing the aggregation of Nodes 1 and 2, column "NN 123" representing the aggregation of Nodes 1, 2 and 3, and so on. The last column "C-TBM" shows the results for the centralized transferable belief model.

Next, we evaluate the scenario using the distributed algorithm. Table 4 shows the edge selection data. Table 5 shows the output of the $\mathcal{R}$-operator. Each column corresponds to the aggregation of two nodes as related to the sequential timing.

Table 2.   BBA assignments for the five nodes.

Set	Node 1	Node 2	Node 3	Node 4	Node 5
∅	0.0	0.0	0.0	0.0	0.0
{a}	0.3	0.0	0.0	0.0	0.2
{b}	0.3	0.4	0.1	0.4	0.4
{c}	0.0	0.4	0.5	0.4	0.1
{a,b}	0.3	0.0	0.0	0.0	0.0
{a,c}	0.0	0.0	0.0	0.0	0.0
{b,c}	0.0	0.0	0.3	0.0	0.0
{a,b,c}	0.1	0.2	0.1	0.2	0.3

Table 3.   Centralized algorithm output with incremental aggregation.

Set	NN 12	NN 123	NN 1234	NN 12345	C-TBM
∅	0.36	0.468	0.5304	0.6334	0.6334
{a}	0.18	0.108	0.0648	0.0451	0.0451
{b}	0.34	0.346	0.3676	0.3070	0.3070
{c}	0.04	0.046	0.0308	0.0125	0.0125
{a,b}	0.06	0.024	0.0048	0.0014	0.0014
{a,c}	0.0	0.0	0.0	0.0	0.0
{b,c}	0.0	0.006	0.0012	0.0004	0.0004
{a,b,c}	0.20	0.002	0.0004	0.0001	0.0001

Table 4.   Temporal edge selection.

Time	t=1	t=2	t=3	t=4	t=5	t=6	t=7
Edge	e_{12}	e_{23}	e_{34}	e_{45}	e_{34}	e_{23}	e_{12}

After each exchange of knowledge between two nodes, the $\mathcal{R}$-operator reveals that the two nodes have the same knowledge value as the operator output. For example, at $t = 5$, the nodes have the same value:

$$\{0.6334, 0.0451, 0.3070, 0.0125, 0.0014, 0.0, 0.0004, 0.0001\}$$

and are consistent with the centralized transferable belief model outputs. Note that the algorithm terminates when the nature of the edge selection process is known [6]. The results demonstrate that the decentralized approach (Table 5) yields the same values as the centralized approach (Table 3).

7.   Conclusions

The decentralized extension of the transferable belief model enables the application of evidence theory to data fusion in critical infrastructure applications.

Table 5.   Distributed algorithm output.

Set	$s_1 \oplus s_2$	$s_2 \oplus s_3$	$s_3 \oplus s_4$	$s_4 \oplus s_5$	$s_3 \oplus s_4$	$s_2 \oplus s_3$	$s_1 \oplus s_2$
$\emptyset$	0.36	0.468	0.5304	0.6334	0.6334	0.6334	0.6334
{a}	0.18	0.108	0.0648	0.0451	0.0451	0.0451	0.0451
{b}	0.34	0.346	0.3676	0.3070	0.3070	0.3070	0.3070
{c}	0.04	0.046	0.0308	0.0125	0.0125	0.0125	0.0125
{a,b}	0.06	0.024	0.0048	0.0014	0.0014	0.0014	0.0014
{a,c}	0.0	0.0	0.0	0.0	0.0	0.0	0.0
{b,c}	0.0	0.006	0.0012	0.0004	0.0004	0.0004	0.0004
{a,b,c}	0.20	0.002	0.0004	0.0001	0.0001	0.0001	0.0001

This is important because centralized computation is ill-suited to critical infrastructure applications due to the associated interdependencies. The case study shows that the decentralized approach produces the same results as the centralized approach, and also demonstrates the utility of fusing data in a distributed manner for interdependent critical infrastructures.

Integrating situational awareness with distributed monitoring is an appealing concept for networked critical infrastructures. This is important because the ability to leverage and share sensor data can significantly enhance system resilience and robustness. For example, in the case of a cyber attack, data from intrusion detection systems coupled with data from standard field sensors can be combined to obtain more accurate belief assessments.

The decentralized approach offers the same advantages as the traditional transferable belief model in terms of its ability to deal with uncertainty. It is important to note, however, that the same disadvantages exist (e.g., exponential growth in the computational complexity with respect to the number of hypotheses). Nevertheless, both approaches are useful for modeling beliefs in interdependent infrastructures for the purpose of enhancing situational awareness.

Acknowledgement

References

[1] S. Boyd, A. Ghosh, B. Prabhakar and D. Shah, Randomized gossip algorithms, *IEEE Transactions on Information Theory*, vol. 52(6), pp. 2508–2530, 2006.

[2] Y. Dalal, A distributed algorithm for constructing minimal spanning trees, *IEEE Transactions on Software Engineering*, vol. 13(3), pp. 398–405, 1987.

[3] S. Das, *High-Level Data Fusion*, Artech House, Norwood, Massachusetts, 2008.

[4] A. Dempster, Upper and lower probabilities induced by a multivalued mapping, in *Classic Works of the Dempster-Shafer Theory of Belief Functions, Studies in Fuzziness and Soft Computing*, R. Yager and L. Liu (Eds.), Springer, Berlin-Heidelberg, Germany, pp. 57–72, 2008.

[5] S. De Porcellinis, S. Panzieri and R. Setola, Modeling critical infrastructure via a mixed holistic reductionistic approach, *International Journal of Critical Infrastructures*, vol. 5(1/2), pp. 86–99, 2009.

[6] A. Gasparri, F. Fiorini, M. DiRocco and S. Panzieri, A networked transferable belief model approach for distributed data aggregation, *IEEE Transactions on Systems, Man and Cybernetics, Part B*, vol. 42(2), pp. 391–405, 2012.

[7] J. Hall and J. Lawry, Generation, combination and extension of random set approximations to coherent lower and upper probabilities, *Reliability Engineering and System Safety*, vol. 85(1-3), pp. 89–101, 2004.

[8] F. Jensen and T. Nielsen, *Bayesian Networks and Decision Graphs*, Springer, New York, 2007.

[9] K. Ng and B. Abramson, Uncertainty management in expert systems, *IEEE Expert*, vol. 5(2), pp. 29–48, 1990.

[10] G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, New Jersey, 1976.

[11] P. Smets and R. Kennes, The transferable belief network, *Artificial Intelligence*, vol. 66(2), pp. 191–234, 1994.

[12] A. Veremme, D. Dupont, E. Lefevre and D. Mercier, Belief assignment on compound hypotheses within the framework of the transferable belief model, *Proceedings of the Twelfth International Conference on Information Fusion*, pp. 498–505, 2009.

[13] F. Voorbraak, A computationally efficient approximation of Dempster-Shafer Theory, *International Journal of Man-Machine Studies*, vol. 30(5), pp. 525–536, 1989.