

# Distributed Belief Propagation for Situation Assessment in Net-Centric Environment

Subrata Das

**Abstract**— A distributed net-centric environment consist of a large variety of data fusion nodes, where each node represents a sensor, software program, machine, human operator, warfighter, or a combat unit. Fusion nodes can be conceptualized as intelligent autonomous agents that communicate, coordinate, and cooperate with each other in order to improve their local situational awareness (SA), and to assess the situation of the operational environment as a whole. In this paper, we describe how we model this net-centric SA problem using a distributed belief propagation paradigm. A local fusion node maintains the joint state of the set of variables modeling a local SA task at hand using Bayesian network (BN) fragments. Local fusion nodes communicate their beliefs and coordinate with each other to update their local estimates of the situation and contribute to the global SA of the environment. We have implemented the propagation paradigm to determine threat out of terrorist dirty bombs with agents searching unstructured intelligence reports for evidence and assessing local situations via BN fragments. The paradigm provides an important foundation of our company's cutting-edge predictive analytics platforms offering to solve enterprise distributed big data search and analytics problems.

**Index Terms**—Net-centric environment; high-level fusion; situation assessment; distributed belief propagation; Bayesian networks.

## I. INTRODUCTION

The concept of distributed fusion [12] refers to decentralized processing environments consisting of autonomous sensor nodes, and additional processing nodes without sensors, if necessary, to facilitate message communication, data storage, relaying, information aggregation, and assets scheduling. Some of the advantages of distributed fusion are reduced communication bandwidth, distribution of processing load, aggregation of distributed and proprietary knowledge sources, and improved system survivability from a single point failure. The distributed fusion concept naturally fits within the net-centric multi-agent paradigm[17]. In this paper, we describe a framework for distributed high-level fusion developed by the author and presented earlier in chapter 10 of [12]. The main development since then is a robust implementation of the framework, making use of both structured and unstructured data sources stored in traditional relational databases and cloud.

This work was supported in part by the U.S. Department of Defense under the Army Grant W15P7T-13-C-A752.

Subrata Das is with Machine Analytics, Cambridge, MA 02138, USA (e-mail: sdas@machineanalytics.com).

As a concrete example of distributed fusion, consider the decentralized processing environment as shown in Figure 1 (left). In this example, we assume there is a high-value target within a region of interest, and that the designated areas A and B surrounding the target are considered to be the most vulnerable. These two areas must be under surveillance in order to detect any probing activities, which indicate a possible attack threat. The sensor coverage in areas A and B, shown in grey, is by an infrared sensor (MASINT) and a video camera (IMINT), respectively. In addition, a human observer (HUMINT) is watching the area in common between A and B. There are two local fusion centers for the two areas to detect any probing activity. The infrared sensor has wireless connectivity with the local fusion center for area A, whereas the video camera has wired connectivity with the local fusion center for area B for streaming video. Moreover, the human observer communicates wirelessly with both local fusion centers. Each of the two local centers fuses the sensor data it receives in order to identify any possible probing activity. The centers then pass their assessments (i.e., higher-level abstraction, rather than raw sensor information, thus saving bandwidth) to another fusion center that assesses the overall threat level, based on the reports of probing activities and other relevant prior contextual information.

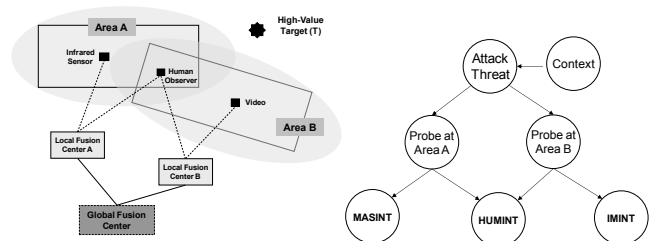


Figure 1: (left) An example distributed fusion environment; (right) A centralized BN model for situation assessment

In a centralized fusion environment, where observations from IMINT, HUMINT, and MASINT are gathered in one place and fused, a BN model, such as the one in Figure 1 (right), can be used for an overall SA. This model handles dependence among sensors and fusion centers via their representation in nodes and interrelationships. A probing activity at an area will be observed by those sensors covering the area, and the lower half of the BN models this. For example, MASINT and HUMINT reports will be generated due to a probing activity at area A. Similarly, IMINT and HUMINT reports will be generated due to a probing activity at area B. The upper half of the BN models the threat of an

attack based on the probing activities at areas A and B, together with other contextual information.

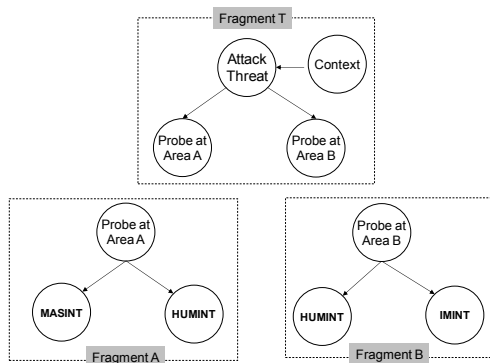


Figure 2: Distributed parts of the BN models

In a decentralized environment [27], as illustrated in Figure 2, each of the three fusion centers contains only a fragment of the above BN model. Local fusion centers A and B assess probing activities based on their local model fragments, and send their assessments to the global fusion center via messages. The global fusion center then uses its own models to determine the overall attack threat. If the same HUMINT report is received by both local fusion centers, the process has to ensure that this common information is used only once; otherwise, there will be a higher-than-actual level of support for a threat to be determined by the global fusion model. This is called the *data incest problem* in a distributed fusion environment, which is the result of repeated use of identical information. *Pedigree* needs to be traced, not only to identify common information, but also to assign appropriate trust and confidence to data sources. An information graph [18], for example, allows common prior information to be found.

For situation and threat assessment in a distributed net-centric environment, each node is an agent representing [3][6][8] a sensor, software program, machine, human operator, warfighter, or a unit. A fusion node maintains the joint state of the set of variables modeling a local SA task at hand. Informally, the set of variables maintained by a fusion node is a *clique* (maximal sets of variables that are all pairwise linked), and the set of cliques in the environment together form a clique network to be transformed into a *junction tree*, where the nodes are the cliques. Thus the cliques of a junction tree are maintained by local fusion nodes within the environment. Local fusion nodes communicate and coordinate with each other to improve their local estimates of the situation, avoiding the repeated use of identical information.

A junction tree can also be obtained by transforming [16] a Bayesian Belief Network (BN) [24][16][7] model representing a global SA model in the context of a mission, thereby contributing to the development of a Common Tactical Picture (CTP) of the mission via shared awareness. Each clique is maintained by a local fusion node. Inference on such a BN model for SA relies on evidence from individual local fusion nodes. We make use of the message-passing inference algorithm for junction trees that naturally fits within distributed NCW environments. A BN structure with nodes and links is a natural fit for distributing tasks in a NCW

environment at various levels of abstraction and hierarchy. BNs have been applied extensively for centralized fusion (e.g., [28][4][21][26]) where domain variables are represented by nodes.

## II. RELATED WORK

There are approaches along these lines, namely Distributed Perception Networks (DPN) [23] and Multiply Section Bayesian Networks (MSBN) [29], but the proposed approach leverages existing algorithms and reduces the overall message flow to save bandwidth. Please refer to [22] for a more detailed account of a junction tree-based distributed fusion algorithm along the lines of the one presented here. The algorithm presented later in the paper, in addition, optimizes the choice of junction tree to minimize the communication and computation required by inference.

There is an abundance of work in the area of distributed agent-based target tracking and, more generally, in the area of distributed fusion. In general, a distributed processing architecture for estimation and fusion consists of multiple processing agents. Here we mention only some of them.

Horling et al. [13] developed an approach to real-time distributed tracking, where the environment is partitioned into sectors to reduce the level of potential interaction between agents. Hughes and Lewis [14] investigated the Track-Before-Detect (identify tracks before applying thresholds) problem using multiple intelligent software agents. Martin and Chang [19] developed a tree based distributed data fusion method for ad hoc networks, where a collection of agents share and fuse data in an ad hoc manner for estimation and decision making.

Graphical models and graphical Bayesian belief networks have been applied extensively by the fusion community to perform estimation [2] and situation assessment [4]. A network structure, modeling a situation assessment problem, with nodes and links is a natural fit for distributing tasks at various levels of abstraction and hierarchy, where nodes represent agents with message flows between agents along the links. An approach along these lines has been adopted by Pavlin et al. [23]. Mastrogianni et al. [20] developed a framework for collaborating agents for distributed knowledge representation and data fusion based on the idea of an ecosystem of interacting artificial entities. Mobile agents [5] have also been employed for distributed fusion.

Mobile agents are able to travel between nodes of a network in order to make use of resources that are not locally available. Mobile agents enable the execution code to be moved to the data sites, thus save the network bandwidth and provide an effective way to overcome network latency. Qi et al. [25] developed an infrastructure for Mobile-agent-based Distributed Sensor Networks (MADSNs) for multisensor data fusion. Bai et al. [1] developed a Mobile Agent-Based Distributed Fusion (MADFUSION) system for decision making in Level 2 Fusion. The system environment consists of a peer-to-peer ad-hoc network in which information may be dynamically distributed and collected via publish/subscribe functionality. Jameson's Grapevine architecture [15] for data fusion integrates intelligent agent technology, where an agent generates the information needs of the peer platform it represents. Gerken et al. [11] embedded intelligent agents into

the Mobile Commander's Associate (MCA) decision aiding system to improve the situational awareness of the commander by monitoring and alerting based on the information gathered.

### III. IMPLEMENTATION

Our approach to complex analytics and fusion<sup>1</sup> is to make use of a computational model and its mobile agent-based distributed belief propagation presented in this paper. Figure 3 presents a Bayesian network model to help in assessing the level of a dirty-bomb threat from a rogue nation. In our model, the site maintaining the root node (for example) continually updates the state (i.e., the probability distribution) of an overall threat based on evidence it receives from its child node, representing terrorism indication and warning, which in turn receives evidence of indications and warnings from its four children, namely, planning, acquisition, making, and deployment. The state of these nodes can be maintained by various sites based on the evidence received from their children nodes. This hierarchical breakdown process continues, and model fragments are determined. A fragment is defined here as a connected sub-network of the entire belief network model. For the purpose of our demonstration we assume all the fragments are two levels deep as shown in Figure 3.

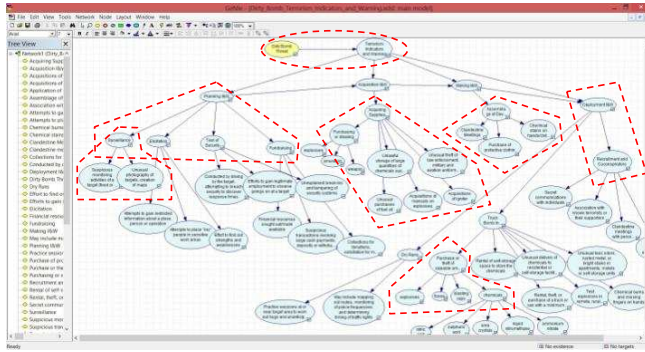


Figure 3: Some fragmented models distributed across remote sites

The specific strategy for evaluating a fragment at a remote site is determined based on the site's capability of accumulating evidence from multiple sources. Figure 4 shows the current state of the interface for controlling and monitoring the distributed execution, with nine list components for the following nine purposes respectively:

1. Text area where the analyst poses a full or partial analytics query in key words.
2. Lists all the model fragments stored in a directory such as the ones from the BN in Figure 3, filters the models based on the analyst query, and lets user select one.
3. The selected dependent model fragments based on the user selection that are to be distributed and maintained across remote sites.
4. Publishes available http addresses of the remote sites running Aglets servers to host computation.

<sup>1</sup> Analytics and data fusion are two sides of the same coin (Das, 2014)

5. Lists search nodes of the fragments (same as the list of model fragments above).
6. Provides probability distributions corresponding to the search nodes as model fragments.
7. Graphically display a selected fragment from the library.
8. Displays messages that are received and also the evidence that are found.
9. Overall assessment of threat which is 0.25 based on evidence searched so far.

Users can dispatch fragments individually by selecting a fragment from the area marked 3 to a remote site selected from the area marked 4 by pressing the Dispatch Agent button. Users can also dispatch all fragments at once just by pressing the Random Dispatch button. Evidence on a child node at a remote site can be set by selecting the node in the area marked 6 and then by pressing the button Set Evidence. Various messages will be passed among fragments as described earlier in the section on complex analytics. These messages will be displayed in the area marked 8. The probability distributions of each child node will be updated in areas marked by 6. To start execution of the analytics model, a user dispatches all fragments at once by pressing the Random Dispatch button in the analytics interface.

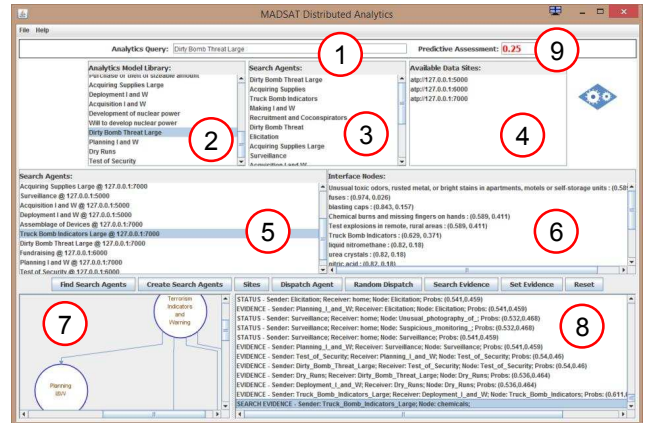


Figure 4: Distributed analytics interface

### IV. DISTRIBUTED FUSION ENVIRONMENTS

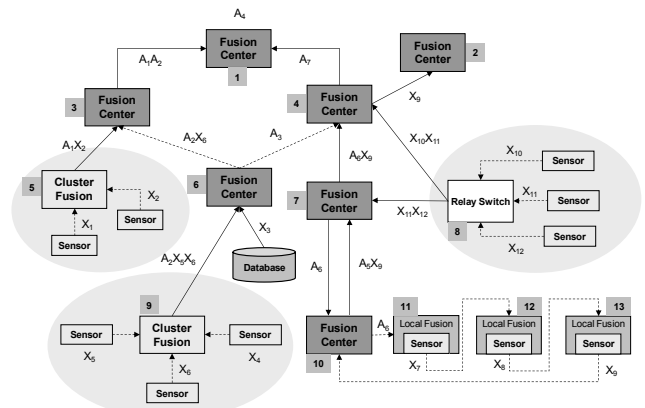


Figure 5: A typical distributed fusion environment

As shown in Figure 5, a typical distributed fusion environment is likely to contain a variety of fusion nodes that do a variety of tasks:

- Process observations generated from a cluster of heterogeneous sensors (e.g., the local fusion centers A and B in Figure 1, and nodes labeled 5 and 9 in Figure 5).
- Process observations generated from a single sensor (e.g., nodes labeled 11, 12, and 13 in Figure 5).
- Perform a task (e.g., Situation Assessment (SA) and Threat Assessment (TA), Course of Action (COA) generation, planning and scheduling, Common Tactical Picture (CTP) generation, collection management) based on information received from other sensors in the environment and from other information stored in databases (e.g., nodes labeled 1, 2, 3, 4, 6, 7, and 10 in Figure 5).
- Relay observations generated from sensors to other nodes (e.g., the node labeled 8 in Figure 5).

As shown in Figure 5, a fusion node receives values of some variables obtained either from sensor observations (X variables) or via information aggregation by other nodes (A variables). Such values can also be obtained from databases. For example, the fusion center labeled 6 receives values of the variables A2, X5, and X6 from the cluster fusion node labeled 9, and values of the variable X3 from a database. Note that an arrow between two nodes indicates the flow of information in the direction of the arrow as opposed to a communication link. The existence of an arrow indicates the presence of at least a one-way communication link, though not necessarily a direct link, via some communication network route. For example, there is a one-way communication link from node 3 to node 1. A reverse communication link between these two nodes will be necessary in implementing our message-passing distributed fusion algorithm to be presented later.

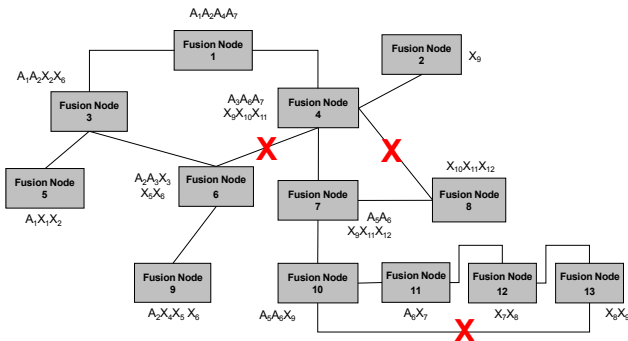


Figure 6: Network of distributed fusion nodes

Each node (fusion center, cluster fusion, relay switch, or local fusion) in a distributed fusion environment has knowledge of the states of some variables, called *local variables*, as shown in Figure 6 (ignore red cross for now). For example, the fusion node labeled 6 has knowledge of the X variables X3, X5, and X6, and A variables A2 and A3. The node receives values of the variables A2, X5, and X6 from the node labeled 9, and the variable X3 from a database. The node

generates values of the variable A3 via some information aggregation operation. On the other hand, fusion node 9 receives measurements X4, X5, and X6 from a cluster of sensors and generates A2; fusion node 8 relays values of the variables X10, X1, and X12 to other nodes; and fusion node 12 obtains measurements of X8 from a single sensor.

There are four possible distributed fusion environments: centralized, hierarchical, peer-to-peer, and grid-based. In a *centralized* environment, only the sensors are distributed, sending their observations to a centralized fusion node. The centralized node combines the sensor information to perform tracking or SA. In a *hierarchical* environment, the fusion nodes are arranged in a hierarchy, with the higher-level nodes processing results from the lower-level nodes and possibly providing some feedback. The hierarchical architecture will be natural for applications where situations are assessed with an increasing level of abstraction along a command hierarchy, starting with the tracking of targets at the bottom level. Considerable savings in communication effort can be achieved in a hierarchical fusion environment. In both *peer-to-peer* and *grid-based* distributed environments, every node is capable of communicating with every other node. This internode communication is direct in the case of a peer-to-peer environment, but some form of “publish and subscribe” communication mechanism is required in a grid-based environment.

## V. ALGORITHM FOR DISTRIBUTED BELIEF PROPAGATION

As mentioned in the introduction, there are two ways in which we can accomplish SA in a distributed environment: 1) each local fusion node maintains the state of a set of variables; 2) there is a BN model for global SA.

In the first case, we start with a distributed fusion environment such as the one shown in Figure 5. Our distributed SA framework in this case has four steps: 1) Network formation; 2) Spanning tree formation; 3) Junction tree formation; and 4) Message passing. The nodes of the sensor network first organize themselves into a network of fusion nodes, similar to the one shown in Figure 6. Each fusion node has partial knowledge of the whole environment. This network is then transformed into a *spanning tree* (a spanning tree of a connected, undirected graph, such as the one in Figure 6, is a tree composed of all the vertices and some or all of the edges of the graph), so that neighbor nodes establish high-quality connections. In addition, the spanning tree formation algorithm optimizes the communication required by inference in junction trees. The algorithm can recover from communication and node failures by regenerating the spanning tree. Figure 6 with (red crosses indicating link severed) describes a spanning tree obtained from the network in Figure 5. The decision to sever the link between nodes 4 and 6, as opposed to between nodes 3 and 6, can be mitigated using the communication bandwidth and reliability information in the cycle of nodes 1, 3, 6, and 4.

Using pairwise communication-link information sent between neighbors in a spanning tree, the nodes compute the information necessary to transform the spanning tree into a junction tree for the inference problem. Finally, the inference problem is solved via message-passing on the junction tree.

During the formation of a spanning tree, each node chooses a set of neighbors, so that the nodes form a spanning tree where adjacent nodes have high-quality communication links. Each node's clique is then determined as follows. If  $i$  is a node and  $j$  is a neighbor of  $i$ , then the variables reachable to  $j$  from  $i$ ,  $R_{ij}$ , are defined recursively as

$$R_{ij} = D_i \cup \bigcup_{k \in \text{nbr}(i) - \{j\}} R_{ki}$$

where  $D_i$  is the set of local variables of node  $i$ . A base case corresponds to a leaf node, which is simply a collection of a node's local variables. If a node has two sets of reachable variables to two of its neighbors that both include some variable  $V$ , then the node must also carry  $V$  to satisfy the running intersection property of a junction tree. Formally, node  $i$  computes its clique  $C_i$  as  $C_i = D_i \cup \bigcup_{\substack{j, k \in \text{nbr}(i) \\ j \neq k}} R_{ji} \cap R_{ki}$

A node  $i$  can also compute its *separator set*  $S_{ij} = C_i \cap C_j$  with its neighbor  $j$  using reachable variables as  $S_{ij} = C_i \cap R_{ji}$ .

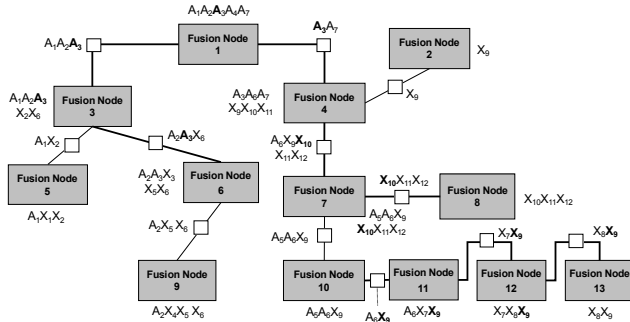


Figure 7: A junction tree from the distributed fusion environment

Figure 7 shows the junction tree obtained from the spanning tree in Figure 6. The variables reachable to a leaf node, for example, fusion node 9, are its local variables  $A_2, X_4, X_5, X_6$ . The variables reachable to an intermediate node, for example, fusion node 1, from its neighboring nodes 3 and 4 are

$$R_{31} = \{A_1, A_2, A_3, X_1, X_2, X_3, X_4, X_5, X_6\}$$

$$R_{41} = \{A_3, A_5, A_6, A_7, X_7, X_8, X_9, X_{10}, X_{11}, X_{12}\}$$

The local variable of the fusion node 1 is  $D_1 = \{A_1, A_2, A_4, A_7\}$ . Therefore, its clique is  $C_1 = \{A_1, A_2, A_3, A_4, A_7\}$ . The formation of a suitable junction tree from a BN model for SA is the only part of our distributed fusion approach that is global in nature.

#### A. Junction Tree Construction and Inference

The *moral graph* of a BN is obtained by adding a link between any pair of variables with a common “child,” and dropping the directions of the original links in the BN. An undirected graph is *triangulated* if any cycle of length greater than 3 has a chord, that is, an edge joining two nonconsecutive nodes along the cycle. The nodes of a junction tree for a graph are the *cliques* in the graph (maximal sets of variables that are all pairwise linked).

Once we have formed a junction tree from either of the above two cases, such as the one in Figure 7, a message-passing algorithm then computes prior beliefs of the variables in the network via an initialization of the junction tree structure, followed by evidence propagation and marginalization. The algorithm can be run asynchronously on each node responding to changes in other nodes' states. Each time a node  $i$  receives a new separator variables message from a neighbor  $j$ , it recomputes its own clique and separator variables messages to all neighbors except  $j$ , and transmits them if they have changed from their previous values. Here we briefly discuss the algorithm, and how to handle evidence by computing the posterior beliefs of the variables in the network.

A junction tree maintains a joint probability distribution at each node, cluster, or separator set in terms of a belief *potential*, which is a function that maps each instantiation of the set of variables in the node into a real number. The belief potential of a set of variables  $\mathbf{X}$  will be denoted as  $\varphi_{\mathbf{X}}$ , and  $\varphi_{\mathbf{X}}(x)$  is the number onto which the belief potential maps  $x$ . The probability distribution of a set of variables  $\mathbf{X}$  is just the special case of a potential whose elements add up to 1. In other words,

$$\sum_{x \in \mathbf{X}} \varphi_{\mathbf{X}}(x) = \sum_{x \in \mathbf{X}} p(x) = 1$$

The marginalization and multiplication operations on potentials are defined in a manner similar to the same operations on probability distributions.

Belief potentials encode the joint distribution  $p(\mathbf{X})$  of the BN according to the following:

$$p(\mathbf{X}) = \frac{\prod_i \phi_{C_i}}{\prod_j \phi_{S_j}}$$

where  $\phi_{C_i}$  and  $\phi_{S_j}$  are the cluster and separator set potentials, respectively. We have the following joint distribution for the junction tree in Figure 7:

$$p(A_1, \dots, A_9, X_1, \dots, X_{12}) = \frac{\phi_{C_1} \phi_{C_2} \dots \phi_{C_{13}}}{\phi_{S_{13}} \phi_{S_{14}} \phi_{S_{24}} \phi_{S_{35}} \dots \phi_{S_{12,13}}}$$

where  $C_i$  represents the variable in clique  $i$  and  $S_{ij} = C_i \cap C_j$  represents the separator set between nodes  $i$  and  $j$ . It is imperative that a cluster potential agrees with its neighboring separator sets on the variables in common, up to marginalization. This imperative is formalized by the concept of local consistency. A junction tree is *locally consistent* if, for each cluster  $\mathbf{C}$  and neighboring separator set  $\mathbf{S}$ , the following holds:  $\sum_{\mathbf{CS}} \phi_{\mathbf{C}} = \phi_{\mathbf{S}}$ . To start initialization, for each cluster  $\mathbf{C}$

and separator set  $\mathbf{S}$ , set the following:  $\phi_{\mathbf{C}} \leftarrow 1$ ,  $\phi_{\mathbf{S}} \leftarrow 1$ . Then assign each variable  $X$  to a cluster  $\mathbf{C}$  that contains  $X$  and its parents  $pa(X)$ . Then set the following:  $\phi_{\mathbf{C}} \leftarrow \phi_{\mathbf{C}} p(X | pa(X))$ .

When new evidence on a variable is entered into the tree, it becomes inconsistent and requires a global propagation to make it consistent. The posterior probabilities can be computed via marginalization and normalization from the

global propagation. If evidence on a variable is updated, the tree requires re-initialization. Next, we present initialization, normalization, and marginalization procedures for handling evidence.

As before, to start initialization, for each cluster  $C$  and separator set  $S$ , set the following:  $\phi_C \leftarrow 1$ ,  $\phi_S \leftarrow 1$ . Then assign each variable  $X$  to a cluster  $C$  that contains  $X$  and its parents  $pa(X)$ , and then set the following:  $\phi_C \leftarrow \phi_C p(X | pa(X))$ ,  $\lambda_X \leftarrow 1$ , where  $\lambda_X$  is the likelihood vector for the variable  $X$ . Now, perform the following steps for each piece of evidence on a variable  $X$ :

- Encode the evidence on the variable as a likelihood  $\lambda_X^{new}$ .
- Identify a cluster  $C$  that contains  $X$  (e.g., one containing  $X$  and its parents).
- Update as follows:  $\phi_C \leftarrow \phi_C \frac{\lambda_X^{new}}{\lambda_X}$ ,  $\lambda_X \leftarrow \lambda_X^{new}$

Now perform a global propagation using the two recursive procedures *Collect Evidence* and *Distribute Evidence*. Note that if the belief potential of one cluster  $C$  is modified, then it is sufficient to unmark all clusters and call only *Distribute Evidence*( $C$ ). The potential  $\phi_C$  for each cluster  $C$  is now  $p(C, e)$ , where  $e$  denotes evidence incorporated into the tree.

Now marginalize  $C$  into the variable as  $p(X, e) = \sum_{C \in \{X\}} \phi_C$ .

Compute posterior  $p(X | e)$  as follows:

$$p(X | e) = \frac{p(X, e)}{p(e)} = \frac{p(X, e)}{\sum_X p(X, e)}.$$

To update evidence for each variable  $X$  on which evidence has been obtained, first update its likelihood vector. Then initialize the junction tree by incorporating the observations. Finally, perform global propagation, marginalization, etc.

## VI. CONCLUSIONS

We have presented an agent based approach to distributed belief propagation in net-centric environments. The approach provides the foundation of the company's predictive analytics products. We are currently enhancing the product with agent-based approach distributed semantic search to find evidence to propagate in Bayesian network fragments. We are investigating the best way to make use of any types of local model fragments such as rules, neural networks, and decision trees.

## VII. REFERENCES

- [1] Bai, L., Landis, J., Salerno, J., Hinman, M., and Boulware, D. (2005). "Mobile Agent-Based Distributed Fusion (MADFUSION) System," *Proceeding of the 8<sup>th</sup> International conference on Information Fusion*, Philadelphia.
- [2] Chong, C-Y and Mori, S. (2004). "Graphical models for nonlinear distributed estimation," *Proceedings of the Conference on Information Fusion*, Vol. I, pp. 614–621.
- [3] Das, S. and Grecu, D. (2000). "COGENT: Cognitive agent to amplify human perception and cognition," *Proc. of the 4<sup>th</sup> Int. Conf. on Autonomous Agents*, Barcelona, June.
- [4] Das, S., Grey, R., and Gonsalves, P. (2002a). "Situation assessment via Bayesian Belief Networks," *Proc. of the 5<sup>th</sup> Int. Conference on Information Fusion*, Annapolis, Maryland.
- [5] Das, S., Shuster, K., and Wu, C. (2002b). "ACQUIRE: Agent-based Complex QUery and Information Retrieval Engine," *Proceedings of the 1<sup>st</sup> International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Bologna, Italy.
- [6] Das, S. (2008a). *Foundations of Decision-Making Agents: Logic, Probability, and Modality*, World Scientific/Imperial College Press, Singapore/London.
- [7] Das, S. (2008b). *High-Level Data Fusion*, Artech House, Norwood, MA.
- [8] Das, S. (2010). Agent-based information fusion, Guest Editorial, *Information Fusion*, Elsevier Science, Vol. 11, pp. 216–219.
- [9] Das, S. (2014). "Computational Business Analytics," Chapman & Hall/CRC Data Mining and Knowledge Discovery Series.
- [10] Durrant-Whyte, H. and Stevens, M. (2006). "Data Fusion in Decentralised Sensing Networks," *Australian Centre for Field Robotics*, The University of Sydney NSW 2006, <http://www.acfr.usyd.edu.au>.
- [11] Gerken, P., Jameson, S., Sidharta, B., and Barton, J. (2003). "Improving Army Aviation Situational Awareness with Agent-Based Data Discovery," *American Helicopter Society 59<sup>th</sup> Annual Forum*, Phoenix, Arizona.
- [12] Hall, D., Liggins, M., and Chong, C. (eds) (2012). *Distributed Data Fusion for Network-Centric Operations*, CRC Press.
- [13] Horling, B., Vincent, R., Mailler, R., Shen, J., Becker, R., Rawlins, K., and Lesser, V. (2001). "Distributed sensor network for real time tracking," *Proceedings of the 5<sup>th</sup> International Conference on Autonomous Agents*, Montreal, 417–424.
- [14] Hughes, E. and Lewis, M. (2009). "An Intelligent agent based Track-Before-Detect system applied to a range and velocity ambiguous radar," *Electro Magnetic Remote Sensing Defence Technology Center (EMRS DTC) Technical Conference*.
- [15] Jameson, S. (2001). "Architectures for Distributed Information Fusion to Support Situation Awareness on the Digital Battlefield," *Proc. of the 4<sup>th</sup> Int. Conf. on Data Fusion*, pp. 7–10.
- [16] Jensen, F. V. (2001). *Bayesian Networks and Decision Graphs*. Springer-Verlag, NY.
- [17] Lichtblau, D. E. (2004). "The critical role of intelligent software agents in enabling net-centric command and control," *Command and Control Research and Technology Symposium, The Power of Information Age Concepts and Technologies*, San Diego, CA.
- [18] Liggins, M. E., Chong, C.-Y., Kadar, I., Alford, M. G., Vannicola, V., and Thomopoulos, S. (1997). "Distributed fusion architectures and algorithms for target tracking," *Proceedings of the IEEE*, 85(1):95–107.
- [19] Martin, T. and Chang, K. (2005). "A distributed data fusion approach for mobile ad hoc networks," *Proceedings of the 8<sup>th</sup> Int. Conference on Information Fusion*, pp. 25–28.
- [20] Mastrogianni, F., Sgorbissa A., and Zaccaria, R. (2007). "A Distributed Architecture for Symbolic Data Fusion,"

*Proceedings of the 20<sup>th</sup> International Joint Conference on Artificial Intelligence (IJCAI)*, Hyderabad, India.

- [21] Mirmoeini, F. and Krishnamurthy, V. (2005). "Reconfigurable Bayesian networks for adaptive situation assessment in battlespace," *Proceedings of the IEEE Conference on Networking, Sensing and Control*, pp. 810–815.
- [22] Paskin, M., and Guestrin, C. (2004). "Robust probabilistic inference in distributed systems," *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence (UAI)*, Banff, Canada.
- [23] Pavlin, G., de Oude, P., Maris, M., and Hood, T. (2006). "Distributed perception networks," *Proc. of the International Conference on Multisensor Fusion and Integration for Intelligent Systems*, Heidelberg, Germany.
- [24] Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo, CA.
- [25] Qi, H., Wang, X., Iyengar, S., and Chakrabarty, K. (2001). "Multisensor Data Fusion in Distributed Sensor Networks Using Mobile Agents," *Proceedings of 5<sup>th</sup> International Conference on Information Fusion*, pp. 11–16.
- [26] Su, X., Bai, P., Du, F., and Feng, Y. (2011). "Application of Bayesian Networks in Situation Assessment," *Intelligent Computing and Information Science, Communications in Computer and Information Science, Volume 134*, Springer Berlin Heidelberg.
- [27] Waldock, A. and Nicholson, D. (2007). "Cooperative decentralised data fusion using probability collectives," *Proc. of the 1<sup>st</sup> Int. Work. on Agent Technology for Sensor Networks*.
- [28] Wright, E., et al, T. (2002). "Multi-entity Bayesian networks for situation assessment," *Proceedings of the 5<sup>th</sup> International Conference on Information Fusion*, pp. 804–811.
- [29] Xiang, Y., Poole, D., and Beddoes, M. (1993). "Multiply sectioned Bayesian networks and junction forests for large knowledge based systems," *Computational Intelligence*, 9(2):171–220, 1993.

**Dr. Subrata Das** is the founder of Machine Analytics, a company in the Boston area providing business analytics and data fusion consultancy services and developing customized solutions for clients in government and industry. Subrata is also providing consulting services to several companies. Subrata's technical expertise includes mathematical logics, probabilistic reasoning including Bayesian belief networks, symbolic argumentation, particle filtering, and a broad range of



computational artificial intelligence techniques.

Subrata recently spent two years in Grenoble, France, as the manager of over forty researchers in the document content laboratory at the Xerox European Research Centre. Subrata guided applied analytics research and development in the areas of unstructured data analyses, machine translation, image processing, and decision-making under uncertainty. Subrata was one of the five-members in the high-profile Xerox task force Knowledge Work 2020, alongside colleagues from the Palo Alto Research Center (PARC), to explore a strategic vision of the future of work.

Before joining Xerox, Subrata held the Chief Scientist position at Charles River Analytics in Cambridge, MA, working on projects funded by DARAP, NASA, US Air Force, Army and Navy, ONR and AFRL. In the past, Subrata held research positions at Imperial College and Queen Mary and Westfield College, both part of the University of London. He received his PhD in Computer Science from Heriot-Watt University in Scotland, a Master's in Mathematics from the University of Kolkata, and an M.Tech from the Indian Statistical Institute.

Subrata has published many journal and conference articles. He is the author of the books *Computational Business Analytics*, published by CRC Press/Chapman and Hall, *High-Level Data Fusion*, published by the Artech House, *Foundations of Decision Making Agents: Logic, Modality, and Probability*, published by the World Scientific/Imperial College Press, and *Deductive Databases and Logic Programming*, published by Addison-Wesley. Subrata has also co-authored the book entitled *Safe and Sound: Artificial Intelligence in Hazardous Applications*, published by the MIT Press.

Subrata served as a member of the editorial board of the *Information Fusion* journal, published by Elsevier Science. He has been a regular contributor, a technical committee member, a panel member, and a tutorial lecturer at the International Conference on Information Fusion. Subrata has published many conference and journal articles, and conceived and developed the in-house tools aText, iDAS and RiskAid.