

Adaptive uncertain information fusion to enhance plan selection in BDI agent systems

Sarah Calderwood, Kim Bauters, Weiru Liu and Jun Hong
School of Electronics, Electrical Engineering and Computer Science,
Queen's University Belfast, United Kingdom
{scalderwood02, k.bauters, w.liu, j.hong}@qub.ac.uk

Abstract—Correctly modelling and reasoning with uncertain information from heterogeneous sources in large-scale systems is critical when the reliability is unknown and we still want to derive adequate conclusions. To this end, context-dependent merging strategies have been proposed in the literature. In this paper we investigate how one such context-dependent merging strategy (originally defined for possibility theory), called largely partially maximal consistent subsets (LPMCS), can be adapted to Dempster-Shafer (DS) theory. We identify those measures for the degree of uncertainty and internal conflict that are available in DS theory and show how they can be used for guiding LPMCS merging. A simplified real-world power distribution scenario illustrates our framework. We also briefly discuss how our approach can be incorporated into a multi-agent programming language, thus leading to better plan selection and decision making.

Keywords—Dempster-Shafer theory; information fusion; context-dependent merging; BDI plan selection.

I. INTRODUCTION

SCADA (Supervisory Control and Data Acquisition) systems have proven to be a powerful and successful technology, deployed in various environments such as power generation and distribution [2]. Developing and managing such systems is complex, but they can be modelled using multi-agent systems [12] such as the BDI framework. Each agent in the BDI framework is modelled by its (B)eliefs (its current belief state), (D)esires (what it wants to achieve) and (I)ntentions (desires the agent is acting upon). However, a BDI agent cannot deal with uncertain, incomplete or conflicting information collected from multiple sources (e.g. sensors or experts). This is often problematic in realistic SCADA settings as the environment is pervaded by uncertainty.

In this paper, we address these issues by extending the BDI framework with a context-dependent merging strategy, originally proposed in the setting of possibility theory [6]. We adapt the work to the setting of Dempster-Shafer (DS) theory, which is a well-understood formal framework for combining multiple sources of evidence. Notably, the explicit representation of conflict and ignorance makes it an ideal model for such complex systems. The ability to represent and reason about the uncertain environment leads to better plan selection strategies. The best plan is selected by determining the largely partially maximal consistent subset (LPMCS) merge of various sources, transforming the merge result into a probability distribution and then using this distribution for decision making. To integrate these components into the BDI

architecture, we show how we need to extend concepts such as a the belief base, as well as provide mechanisms for conditional merging to alleviate the computational overhead associated with context-dependent merging.

This paper is organised as follows. Section II provides some necessary preliminaries, while Section III presents a power distribution scenario as a running example. In Section IV we adapt the LPMCS merge from possibility theory to the setting of DS theory and we show in Section V how the ability to model and reason about uncertain and conflicting information leads to better plan selection in a BDI context. Finally in Section VI we conclude the paper and discuss related work.

II. PRELIMINARIES

Dempster-Shafer [11] is a theory of evidence that is well-suited for dealing with epistemic uncertainty and sensor fusion. The *frame of discernment* $\Omega = \{\omega_1, \dots, \omega_n\}$ is defined as the set of possible events, one of which is true at a particular time. The *basic belief assignment* (*bba*), or *mass function*, is a mapping $m : 2^\Omega \rightarrow [0, 1]$ if $m(\emptyset) = 0$ and $\sum_{A \subseteq \Omega} m(A) = 1$. Intuitively, $m(A)$ defines the proportion of evidence that supports A , but none of its specific subsets. Whenever $m(A) < 0$ for $A \subseteq \Omega$ it is called a *focal element* of m .

One of the best known rules to merge bbas is *Dempster's rule of combination*, or *conjunctive rule*, which is defined as:

$$m(\emptyset) = 0, m(A) = (m_i \otimes m_j)(A) = c \sum_{B \cap C = A \neq \emptyset} m_i(B) m_j(C)$$

with c a normalisation constant, given by $c = \frac{1}{1-K}$ with $K = \sum_{B \cap C = \emptyset} m_i(B) m_j(C)$. The effect of the normalisation constant c , with K the degree of conflict between m_i and m_j , is to redistribute the mass value assigned to the empty set. As such, Dempster's rule is only well-suited to merge sources with a low degree of conflict. Dubois and Prade's *disjunctive consensus rule* [5], on the other hand, is defined as:

$$m(A) = (m_i \oplus m_j)(A) = \sum_{B \cup C = A} m_i(B) m_j(C)$$

Notably, the disjunctive rule omits normalisation and incorporates all conflict. Throughout this paper, irrespective of the merging rule, we will denote the result of merging m_i with m_j as m_{ij} . Also, we will interchangeably use evidence sources and bbas when working with individual sources (e.g. a sensor) to emphasise the underlying intuition. We assume that these

TABLE I
EVIDENCE OBTAINED AT TIME δ WITH DISCOUNTING FACTOR α

evidence	source	value	time δ	hypotheses	discount α
E_1	SR ¹	49.8 (± 0.2)	δ	(n)	0.95
E_2	SR	49.7 (± 0.1)	δ	(n)	0.90
E_3	SR	49.4 (± 0.1)	δ	(n, l)	0.56
E_4	SR	49.7 (± 0.2)	$\delta(-5)$	(n)	0.75
E_5	EE ²	(50.0 - 50.4)	δ	n	0.50
E_6	EE	(49.9 - 50.2)	δ	n	0.85

¹ Sensor Reading ² Expert Estimation

evidence sources have been converted into bbas as needed, e.g. in $E_1 \oplus E_2$ with E_1 and E_2 both evidence sources.

When dealing with multiple sources, their frames of discernment can be distinct from each other. A multi-valued mapping [9] one frame of discernment Ω_1 , to another frame Ω_2 where $\Gamma : \Omega_1 \rightarrow 2^{\Omega_2}$. In other words, a single element from Ω_1 is mapped to a subset of Ω_2 . To reflect the reliability of evidence we apply a discounting factor to each bba. Shafer's discounting technique [3] is defined as:

$$m_\alpha(A) = \alpha \cdot m(A), \forall A \neq \Omega, m_\alpha(\Omega) = (1 - \alpha) + \alpha \cdot m(\Omega)$$

where $\alpha = 0$ represents a fully reliable source and $\alpha = 1$ an unreliable source.

To measure the (external) conflict between two bbas we use the distance measure proposed by Jousselme [7], defined for two bbas m_i and m_j on Ω as:

$$d(m_i, m_j) = \sqrt{\frac{1}{2}(\vec{m}_i - \vec{m}_j)^T D(\vec{m}_i - \vec{m}_j)}$$

with \vec{m}_i and \vec{m}_j the vector representations of m_i and m_j . We have that \vec{m}^T is the transpose of vector \vec{m} and D is a $2^\Omega \times 2^\Omega$ similarity matrix whose elements are $D(A, B) = |A \cap B| / |A \cup B|$, with A and B subsets of Ω .

Finally, Smet's pignistic model [3] allows decisions to be made on individual hypotheses by transforming a bba m on Ω into a probability distribution $BetP_m$ over Ω such that:

$$BetP_m(\omega) = \sum_{A \subseteq \Omega, \omega \in A} \frac{m(A)}{|A|}$$

with $|A|$ the number of elements in subset A .

III. POWER DISTRIBUTION SCENARIO

Throughout this paper, we use a power distribution scenario as a running example. To monitor the overall health of a power system i.e. to find its true environmental state, a number of sources with various levels of granularity are considered. We use frequency sensors to determine the frequency of oscillations of alternating current (AC) in the power grid, nominally 50.0 Hz and expert estimations of the state (i.e. frequency) of a given grid location at a given time as in Table I.

Sensor readings provide a specific numerical value and are only accurate up to a given range, e.g. a reading of 50.0 Hz implies the frequency is within 49.8 Hz and 50.2 Hz. Sensor readings will map to the hypotheses (l)ow, (n)ormal, (h)igh where $\Gamma(\{49.0, \dots, 49.5\}) = l$; $\Gamma(\{49.5, \dots, 50.5\}) = n$; $\Gamma(\{50.5, \dots, 51.0\}) = h$. Experts will not give an exact value, but will instead only provide their hypotheses. Still, hypotheses among experts may

not be directly comparable. A frequency that is normal for one expert may already be considered as low by another. As such, we allow an overlap of values between hypotheses where $\Gamma(\{49.0, \dots, 49.7\}) = l$; $\Gamma(\{49.5, \dots, 50.5\}) = n$; $\Gamma(\{50.3, \dots, 51.0\}) = h$. Using this scenario, the frame of discernment that the sources have in common is $\Omega = \{l, n, h\}$.

To estimate the basic belief assignment (bba) of each piece of evidence we apply a reliability discounting factor using Shafer's discounting technique to account for sensor accuracy and the time it was obtained. For sensor readings we can assign probability to $\{l, n, h\}$, whereas for expert estimations we can assign probability to the set of all possible subsets of Ω . For example, given E_1 in Table I we assign α to $\{n\}$ and $1 - \alpha$ to $\{l\}$ as sensor accuracy causes the value to become closer to a low hypotheses. Table II shows our estimated bbas after applying all discounting.

TABLE II
BASIC BELIEF ASSIGNMENTS (BBAS) AFTER APPLYING DISCOUNTING

evidence	m_i^{bba}	\emptyset	$\{l\}$	$\{n\}$	$\{h\}$	$\{l, n\}$	$\{l, h\}$	$\{n, h\}$	Ω
E_1	m_1	0	0.05	0.95	0	0	0	0	0
E_2	m_2	0	0.10	0.90	0	0	0	0	0
E_3	m_3	0	0.56	0.44	0	0	0	0	0
E_4	m_4	0	0.25	0.75	0	0	0	0	0
E_5	m_5	0	0	0.50	0	0	0.05	0.30	0.15
E_6	m_6	0	0	0.85	0	0	0.05	0.05	0.05

IV. CONTEXTUAL MERGING IN DS THEORY

Merging uncertain information is often problematic. For example, a well-known limitation of Dempster's rule of combination (i.e. a conjunctive merge) is that it cannot be used to merge sources that are too much in conflict with each other. Similarly, disjunctively merging sources that are too non-specific will result in total ignorance as to the true value.

In this section, we explore ideas presented in [6] in the setting of possibility theory and adapt to DS theory. In particular, we find suitable replacement for the measures used in their work to determine the context, including measures for internal and external conflict as well as measures to identify the quality of any given source without knowing the actual true value. By using DS theory, we have the benefit of a more expressive framework which can more accurately represent the various forms of uncertainty in the environment.

A. Quality of Information

The best source is commonly the source that is the most specific, while exhibiting the least amount of internal conflict. In DS theory we can measure specificity by considering a measure for non-specificity which is defined by [8]:

$$N(m) = \sum_{A \in F} m(A) \log_2 |A|$$

with m a mass function, F a set of focal elements and $|A|$ the cardinality of A . When $N(m) = 0$ a source m is totally specific while $N(m) = \log_2 |\Omega|$ indicates total ignorance.

To measure internal conflict we consider a measure for discord. When evidence is uncertain (or in disagreement), measuring the internal conflict within subsets provides an idea

TABLE III
UNCERTAINTY MEASURES OF THE SCENARIO IN DS THEORY

m_i	$N(m)$	$S(m)$	m_i	$N(m)$	$S(m)$	m_i	$N(m)$	$S(m)$
m_1	0	0.29	m_2	0	0.47	m_3	0	0.99
m_4	0	0.81	m_5	0.59	0.38	m_6	0.18	0.32

of how coherent they are. The measure of strife [8], which improves upon earlier approaches for measuring discord [8], is defined as:

$$S(m) = - \sum_{A \in F} m(A) \log_2 \sum_{B \in F} m(B) \frac{|A \cap B|}{|A|}$$

with m a mass function and F a set of focal elements. We have that $S(m) = 0$ for a source m that has no conflict, while $S(m) = \log_2 |\Omega|$ indicates total conflict.

Using the measures for non-specificity and discord allows us to assess the quality of each source and, thus, allows us to rank sources based on their quality. We have:

Definition 1. (adapted from Def. 5 from [6]) Let E_1 and E_2 be two evidence sources represented by mass functions m_1 and m_2 , respectively. The quality of m_1 is said to be better than that of m_2 , denoted as $m_1 \prec m_2$, if $N(m_1) < N(m_2)$ or $N(m_1) = N(m_2)$ and $S(m_1) < S(m_2)$ holds. Additionally, we say that two sources are indistinguishable, denoted as $m_1 \sim m_2$ when $N(m_1) = N(m_2)$ and $S(m_1) = S(m_2)$. We say that m_1 is better or equal to m_2 , denoted as $m_1 \preceq m_2$, when either $N(m_1) < N(m_2)$ or $m_1 \sim m_2$.

Example 1. Consider the mass functions m_i associated with the evidence source E_i as provided in Table III. Based on the quality of m_i , the evidence sources can be ordered as $m_1 \prec m_2 \prec m_4 \prec m_3 \prec m_6 \prec m_5$, i.e. the best quality source is m_1 and the least quality source is m_5 .

B. Conditions of Conjunctive Merging

The next step is to merge a *largely partially consistent* subset of the sources conjunctively. In particular, we apply certain thresholds to both the internal conflict of the merged result and the external conflict of the different sources to merge to determine if we can further extend this subset. These thresholds are determined by the knowledge engineer, as they are domain specific, and reflect the degree of uncertainty we are willing to tolerate. We have:

Definition 2. (adapted from Def. 7 from [6]) Let E_1 and E_2 be two evidence sources represented by mass functions m_1 and m_2 , respectively. Let m_{12} be the conjunctively merged result of m_1 and m_2 . Let K be the associated degree of conflict with that merge. These mass functions can be merged conjunctively when $K \leq \varepsilon_{\text{conflict}}$ and $S(m_{12}) \leq \varepsilon_{\text{strife}}$ where $\varepsilon_{\text{conflict}} \in [0, 1]$ is a pre-defined threshold of the maximum degree of conflict we tolerate and $\varepsilon_{\text{strife}} \in [0, \log_2 |\Omega|]$ is a maximum degree of strife we are willing to allow. These mass functions should be merged disjunctively, otherwise.

Example 2. Consider the power distribution scenario. Let $\varepsilon_{\text{conflict}} = 0.25$ and $\varepsilon_{\text{strife}} = 0.30$ as the maximum degree of conflict and strife. These thresholds, given the scenario are acceptable to minimise the internal and external conflict within

and between mass functions. For evidence sources E_1 and E_2 we obtain that, given these thresholds, m_1 and m_2 can be merged conjunctively as $K = 0.14$ and $S(m_{12}) = 0.05$. However, for E_1 and E_3 , m_1 and m_3 are not safe to merge conjunctively as $K = 0.55$ and $S(m_{13}) = 0.33$.

Contrary to [6], a measure does not exist in DS theory to measure the degree of information loss to determine if applying a conjunctive rule is more suitable. Identifying and evaluating such a measure is therefore left for future work.

C. Creating a Largely Partially Maximal Consistent Subset

An LPMCS is the result of merging sources that are largely in agreement with each other. We measure the distance between sources to assess how similar they are to a given reference source, so as to maximise the number of sources we merge in a single LPMCS. In [6], quality measures are used to define an implicit distance. In DS theory, distance measures are common to measure conflict. In this paper we use the widely regarded Jousselme distance [7] to accurately measure the similarity between sources.

Definition 3. (adapted from Def. 8 from [6]) Let E_r be a reference source represented by mass function m_r . Let E_1 and E_2 be two evidence sources represented by mass functions m_1 and m_2 , respectively. Let m_{r1} and m_{r2} be the conjunctively merged results of m_r and m_1 , and m_r and m_2 respectively. Let D be a distance measure. The distance of m_1 is as close or equal to that of m_2 with regard to a reference source, denoted as $m_1 \preceq_{m_r} m_2$, when $D(m_r, m_{r1}) \leq D(m_r, m_{r2})$.

Example 3. Consider the power distribution scenario with evidence sources E_i represented by m_i . The reference source is m_1 (from Definition 1). The results of applying the Jousselme distance measure to m_1 and m_{1i} for $i = 2, \dots, 6$ are 0.05, 0.51, 0.20, 0.33 and 0.09, respectively. Since we find m_1 is the closest to m_{12} , we merge m_1 with m_2 .

Using the distance measure, we define a *preferred sequence for merging*. The first source in the sequence is the reference source with the remaining sources ordered given their distance.

Definition 4. (adapted from Def. 9 from [6]) Let E, E_1, \dots, E_k be $k+1$ evidence sources with corresponding mass functions m, m_1, \dots, m_k . Let m be the reference source. The preferred sequence for merging is $(m, m_{i_1}, \dots, m_{i_j}, m_{i_{j+1}}, \dots, m_{i_k})$ with $1 \leq j < k$ and $m_{i_j} \preceq_{m^{j-1}} m_{i_s}$ with $j < s \leq k$, where m^{j-1} is the conjunctively merged result of the first j sources in the sequence that meet the specified thresholds.

Example 4. Consider the power distribution scenario, the preferred sequence for merging is $(m_1, m_2, m_6, m_4, m_5, m_3)$.

We can now define the concept of a largely partially maximal consistent subset (LPMCS).

Definition 5. (adapted from Def. 10 from [6]) Let $m = \{m, m_{i_1}, \dots, m_{i_j}, m_{i_{j+1}}, \dots, m_{i_k}\}$ be the preferred sequence of merging with reference m . A subset $S_m = \{m, m_{i_1}, \dots, m_{i_j}\}$ is called an LPMCS with reference m if the bbas in S_m can be merged conjunctively but bbas in $S_m \cup \{m_{i_{j+1}}\}$ cannot.

TABLE IV
RESULTS OF LPMCS, CONJUNCTIVE AND DISJUNCTIVE MERGE

merge	\emptyset	$\{l\}$	$\{n\}$	$\{h\}$	$\{l, n\}$	$\{l, h\}$	$\{n, h\}$	Ω
LPMCS	0	0	0.44	0	0.56	0	0	0
conjunctive	0	0	1.0	0	0	0	0	0
disjunctive	0	0	0.12	0	0.31	0	0.08	0.49

TABLE V
RESULTS OF PIGNISTIC TRANSFORMATION

	l	n	h
LPMCS	0.28	0.72	0
conjunctive	0	1.0	0
disjunctive	0.32	0.48	0.21

Example 5. Consider Example 4, the corresponding LPMCS are $\{m_1, m_2, m_6, m_4, m_5\}$ and $\{m_3\}$. For the first subset we cannot conjunctively merge m_{12645} with m_3 since we exceed the thresholds for conflict and strife i.e. $K = 0.55$ and $S(m_{126453}) = 0.33$, respectively. The process is then repeated for the remaining sources.

D. Result of LPMCS Merge

Finally, once we cannot merge any more sources conjunctively, we perform a disjunctive merge using Dubois and Prade’s disjunctive consensus rule. The results are shown in Table IV. For comparative purposes we also apply a conjunctive merge and disjunctive merge. While a conjunctive merge provides a similar result in this case, it should be noted that a conjunctive merge would have failed if a trustworthy source becomes unreliable, in which case the disjunctive merge would prove valuable. The LPMCS merge, on the other hand, will provide acceptable results in all these situations. The pignistic transformation of all three merges is shown in Table V.

V. INTEGRATING LPMCS MERGE INTO AGENTSPEAK

We now propose a new framework to integrate AgentSpeak, a multi-agent programming language, with LPMCS merge. In particular, we show how merging different (conflicting) sources of information can help improve plan selection by enabling a suitable plan to be selected and executed based on the agents beliefs about the current environment.

Before we integrate the LPMCS merge into AgentSpeak, we provide a brief overview of AgentSpeak. An AgentSpeak agent \mathbb{A} is represented as a tuple $\langle Bb, Pl, E, A, I, S_\varepsilon, S_\mathcal{O}, S_\mathcal{T} \rangle$ where respectively we have its belief base (a set of beliefs), plan library (a set of plans), event set, action set, intention set and three selection functions [10]. An AgentSpeak agent works in a continuous reasoning cycle where it senses its environment and generates appropriate events as the environment changes. There may be multiple events in the event set E , and selection function S_ε is used to select the event that the agent will deal with in the current reasoning cycle. The plans that are applicable to deal with this event are selected from Pl . Again, there may be many plans to deal with a single event and selection function $S_\mathcal{O}$ selects one of these plans. This plan then turns into an intention that may require multiple reasoning cycles to execute. To ensure a good scheduling of these intentions, selection function $S_\mathcal{T}$ selects the intention

which executes during this cycle (as many cycles may be needed to fully execute any given intention). The cycle then repeats until there are no events left to deal with and there are no more intentions to execute.

In [10], a belief atom is defined as $b(t)$ where b is a n -ary predicate symbol. If $b(t)$ and $c(t)$ are belief atoms, $b(t)$, $c(t)$, $b(t) \wedge c(t)$ and $\neg b(t)$ are beliefs. Achievement and test goals are defined as $!g(t)$ and $?g(t)$ respectively where g is a predicate symbol. Considering the belief atoms and goals, triggering events are defined as $+b(t)$, $-b(t)$, $+!g(t)$, $-!g(t)$, $+?g(t)$, $-?g(t)$. The operators $+$ and $-$ denote addition and deletion of a belief or goal respectively. An action is defined as $a(t)$ where a is an action symbol. In all these cases t_1, \dots, t_n are terms.

A. Collecting and Modelling Evidence

From the continual sensing of the environment, an agent will acquire percepts representing particular properties of the current state of the environment. We have:

Definition 6. Let a percept be represented as $percept(E_i, H)$ where E_i is an evidence source and X is a subset of a frame that represents the observable outcomes from source E_i .

Example 6. Consider the power distribution scenario. From sensing the environment for two sources E_1 and E_6 we have $percept(E_1, \{50.0\})$ and $percept(E_6, \{normal\})$, for a frequency sensor and an expert estimation, respectively.

Each percept collected from the environment revises an agent’s corresponding beliefs to reflect the current state of the environment. In particular, a percept from a frequency sensor will map its value to a hypothesis and generate a bba to model uncertain information, which is included in the agent’s $Bb+$ base. We have:

Definition 7. Let Bb be a belief base and $\Omega = \{\omega_1, \dots, \omega_n\}$ be a frame of discernment. Let a belief atom from a $percept(E_i, H_i)$ be represented in Bb as $senseSource(E_i, S, H_i)$ where E_i is an evidence id, S is the source type and H_i is a subset of Ω .

Example 7. Consider the evidence collected from sources E_1 and E_6 in Table I. In the belief base Bb , they are represented as $senseSource(E_1, SR, \{49.8, \dots, 50.2\})$ and $senseSource(E_6, EE, \{normal\})$, respectively where SR stands for sensor reading, and EE stands for expert estimation.

Note that belief atoms imply that pieces of evidence are certain. For example, consider E_1 and E_6 in Example 7, the bbas are $m_1(\{n\}) = 1$ and $m_6(\{n\}) = 1$ respectively, for the two belief atoms. Since sources are unreliable, a discounting function using the corresponding reliability factor is applied to derive new bbas. For example, for E_1 with reliability factor 0.95, m_1 is derived with $m_1(\{n\})=0.95$, $m_1(\{l\})=0.05$, $m_1(S)=0$ for all other subsets $S \subseteq \{l, n, h\}$.

Since a belief base can only be used to store belief atoms, we need to extend the notion of a belief base so that we can associate with each belief atom a newly derived bba. We have:

Definition 8. Let $b_1 = senseSource(E_1, S, H_1)$ be a belief atom in Bb , and m_{b_1} be a bba derived from b_1 after applying

the discounting function, then (b_1, m_{b_1}) is a belief atom in $Bb+$, which is referred to as an extended belief atom, and $Bb+$ is referred to as an extended belief base.

Example 8. Consider Example 6 and 7. Given $\Omega = \{l, n, h\}$, a belief atom in $Bb+$ is (b_1, m_{b_1}) where $b_1 = \text{senseSource}(E_1, SR, \text{normal})$; $m_{b_1}(\{n\}) = 0.95$; $m_{b_1}(\{l\}) = 0.05$; $m_{b_1}(\emptyset) = m_{b_1}(\{h\}) = m_{b_1}(\{l, n\}) = m_{b_1}(\{l, h\}) = m_{b_1}(\{n, h\}) = m_{b_1}(\Omega) = 0$.

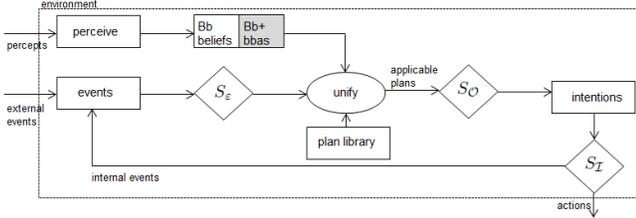


Fig. 1. The reasoning cycle for an extended AgentSpeak agent.

B. Conditions of LPMCS Merge

In AgentSpeak, the LPMCS merge (from Section IV) is executed when a new percept obtained from some source has an associated bba that is considerably in conflict with the current bba that we have in $Bb+$ for that source. To identify the (degree of) conflict, we measure the distance between both bbas. If the new bba is sufficiently similar to the existing bba we assume that the real state has not changed in any meaningful way and we do not perform a new merge. This helps us to avoid the computational cost involved with merging and to ensure the reactivity of the agent (a key property of any BDI agent). Otherwise, a merge needs to be executed to ensure that the beliefs of the agent are up-to-date. We have:

Definition 9. Let E_1, \dots, E_k be k evidence sources with corresponding mass functions m_1, \dots, m_k and $k \geq 2$. Let E_i be any evidence source with an existing mass function m_i and a new mass function m_i^{new} . Let D be a distance measure. The mass functions $m_1, \dots, m_{i-1}, m_i^{new}, m_{i+1}, \dots, m_k$ should be merged using LPMCS merge when $D(m_i, m_i^{new}) \geq \varepsilon_{dist}$ with ε_{dist} a threshold of the similarity between m_i and m_i^{new} .

Example 9. Consider the power distribution scenario. Let $\varepsilon_{dist} = 0.15$ as the maximum degree of conflict allowed between an existing bba and a new bba so that they are sufficiently similar. Let an evidence source E_1 have a previous reading $m_1(\{n\})=0.95$, $m_1(\{l\})=0.05$. For a new reading with $m_1^{new}(\{l\})=0.95$, $m_1^{new}(\{n\})=0.05$, the distance is 0.90. As such the LPMCS merge is executed. On the other hand, if a new reading is $m_1^{new'}(\{n\})=0.85$, $m_1^{new'}(\{l\})=0.15$, the distance is 0.10. An LPMCS merge is not executed.

C. Result of LPMCS Merge

After the LPMCS merge, a new belief atom is inserted into $Bb+$ to reflect the effect of combined evidence. This is represented as (b^n, m^n) , where m^n is the merged bba and b^n is a belief atom in Bb defined as $\text{combinedSource}(N, -, -)$ with N the number of sources that were combined and ‘ $-$ ’ reflecting the Nil source type and Nil hypotheses.

Definition 10. Let P be the pignistic probability function over Ω . For every $\omega_i \in \Omega$, when $P(\omega_i) \geq \varepsilon_{true}$, belief atom $\text{default}(N, -, \omega_i)$ is added to the agent’s belief base Bb , where default is used to denote this belief atom is not from sensor readings, but from a combined result.

Example 10. Consider the power distribution scenario with $\Omega = \{l, n, h\}$ representing the states (l)ow, (n)ormal, (h)igh respectively. Let $\varepsilon_{true} = 0.70$ be the threshold where a state is acceptable. A new belief atom in $Bb+$ is represented as (b^n, m^n) where $b^n = \text{combinedSource}(6, -, -)$ and $m^n(\{n\}) = 0.44$; $m^n(\{l, n\}) = 0.56$. Now, after transforming this bba m^n , we get $P(n) = 0.72$, $P(l) = 0.28$, and $P(h) = 0$. Since $P(n) > \varepsilon_{true}$, a new belief atom $\text{default}(6, -, n)$ is inserted into the agent belief base Bb . Using this example, the grid is in a normal state. However, if no probability of any element in Ω exceeds the threshold, no new belief atoms are generated, so an agent has the same beliefs as prior to this new observation (reflecting that no significant changes occurred in the environment).

In this example the conjunctive merge provides the correct results. Due to space constraints, we do not consider an example where a faulty sensor could cause the conjunctive merge to incorrectly associate complete trust with the incorrect result of that sensor. In such a case, a LPMCS and disjunctive merge would correctly conclude that the actual state of the environment is not known with a degree of probability given that a conclusive result cannot be derived. As such the agent is ignorant about the actual state of the environment resulting in updating the belief and selecting a plan accordingly.

D. Plan Selection

Plan selection is an important component of any agent-based system as it affects how the agent behaves based on its beliefs. In the power distribution setting, the goal of an agent is to select plans to achieve a safe and efficient supply of electrical power to meet consumer demand. This might involve sub-goals such as shutting down a substation when the voltage is outside of the normal range. However, selecting the best plans becomes error-prone when information about the environment is incomplete, uncertain or in conflict. A plan is defined in [10] as $\lambda : b_1 \wedge \dots \wedge b_n \leftarrow h_1; \dots; h_n$ with λ a triggering event, b_1, \dots, b_n the context of the plan, $h_1; \dots; h_n$ the body of a plan. Intuitively, the triggering event determines the events for which the plan can be used, the context determines when a plan is applicable and the plan body is the sequence of actions/subgoals to execute in order to accomplish the goal.

Example 11. Consider the power distribution scenario where agent \mathbb{A} represents a distribution substation. In the belief base of \mathbb{A} , a belief atom such as $\text{default}(N, -, \omega_i)$ records the result from merging evidence where ω_i is either low, normal or high. Other belief atoms include *distributing* and *calibrating* where power is distributing and sensors are calibrated and adjusted. Actions to be executed by \mathbb{A} include *senseSources*, *distribute* and *switchOff* which indicate obtaining sensor readings, the power distributes and the system switches off. Given the environment, the following subset of plans specify the

behaviour of \mathbb{A} . Specifically, plans P_1 to P_3 use merged results, while P_4 uses a single reading. Plan P_1 , is triggered when power is available on the grid in which case the agent proceeds by gathering percepts to continuously validate the state of the grid. In particular, if power is distributing and the grid is normal, the agent continues to distribute power, while sensing for future fluctuations.

```
P1: +!power : default(6,-,normal) & distributing ←
      distribute; senseSources; !power.
```

Plan P_2 , is applicable when the pignistic result is high and power is distributing. The agent will send an alert, switch off and calibrate the sensors.

```
P2: +!power : default(6,-,high) ← alert; calibrate;
      switchOff; !switchOn.
```

Plan P_3 , is triggered when the system is switched on and applicable when the sensors have calibrated and power is not distributing. The agent will then take the actions to distribute power and gather percepts.

```
P3: +!switchOn : calibrated & not distributing ←
      distribute; senseSources; !power.
```

Plan P_4 , is applicable when a single reading is normal and power is distributing. The agent will gather the percept.

```
P4: +!power : default(1,-,normal) & distributing ←
      senseSources; !power.
```

Clearly, the ability of the agent to merge uncertain information has a positive effect on plan selection. In particular, it becomes possible to create plans that are only applicable when sufficient sources agree on the state of the grid, while we retain the ability to base our decisions on only a single source (e.g. when only a single sensor is available).

VI. RELATED WORK AND CONCLUSION

A context-dependent form of merging was suggested in [3] for use in DS theory, where maximally consistent subsets are created from uncertain and conflicting sources. However, their approach only allows to discount all bbas with the same reliability factor, whereas our approach allows a reliability to be associated with each source. Furthermore, our work integrates tightly with BDI, allowing for uncertainty reasoning within a BDI setting. Several approaches to combining BDI with uncertainty modelling have been proposed in the literature. In [4], an agent collects (uncertain) percepts which are fed into a probabilistic graphical model (PGM), revising the agent's epistemic state after uncertainty propagation. Beliefs are updated to a classical belief base where beliefs are either true or false. Further to [4], work in [1] uses the BDI architecture CANPLAN to consider an uncertain belief base where an agent can reason about uncertainty on its own. Actions are also modelled that are affected by uncertainty. Contrary to those approaches, our work focuses on merging different sources of information, which is not considered in the aforementioned papers.

In this paper we presented a largely partially maximal consistent subset (LPMCS) merging strategy in the Dempster-Shafer (DS) setting. In particular, we adapted an existing approach from the setting of possibility theory and show how the

various measures required for an LPMCS merge are available in the DS setting. Furthermore, we integrated and showed this merging strategy in a BDI setting to model uncertainty and reason about merged sensor readings. Specifically, we extended the classical belief base of a BDI agent to accommodate basic belief assignments, introduced a mechanism for conditional merging and showed that the new framework allows for better plan selection through considering the uncertainty of the environment. As for future work, we are going to explore the potential of an iterative LPMCS merge, as well as a full evaluation of our framework. Such a merge would avoid the difficulties associated with a conditional merge (e.g. the context-dependent nature of when such a merge should occur) could further improve the computational characteristic of our new BDI framework.

REFERENCES

- [1] K. Bauters, W. Liu, J.Hong, C. Sierra, and L. Godo. Can(Plan)+: Extending the operational semantics for the BDI architecture to deal with uncertain information. In *Proc. of UAI'14*, 2014.
- [2] S.A. Boyer. SCADA: Supervisory control and data acquisition. In *International Society of Automation*, 2000.
- [3] F. Browne, N. Rooney, W. Liu, D. Bell, H. Wang, P. Taylor, and Y. Jin. Integrating textual analysis and evidential reasoning for decision making in engineering design. In *Knowledge Based Systems*, volume 52, pages 165–175, 2013.
- [4] Y. Chen, J. Hong, W. Liu, L. Godo, C. Sierra, and M. Loughlin. Incorporating PGMs into a BDI architecture. In *Proc. of PRIMA'13*, pages 54–69, 2013.
- [5] D. Dubois and H. Prade. On the combination of evidence in various mathematical frameworks. In *Reliability Data Collection and Analysis*, pages 213–241. Springer Netherlands, 1992.
- [6] A. Hunter and W. Liu. A context-dependent algorithm for merging uncertain information in possibility theory. In *IEEE Transactions on System, Man and Cybernetics*, volume 38(6), pages 1385–1397, 2008.
- [7] A. Jousselme and P. Maupin. Distances in evidence theory: Comprehensive survey and generalizations. In *International Journal of Approximate Reasoning*, volume 53(2), pages 118–145, 2012.
- [8] G.J. Klir and B. Parviz. A note on the measure of discord. In *Proc. of UAI'92*, pages 138–141, 1992.
- [9] W. Liu, J. Hughes, and M. McTear. Representing heuristic knowledge in DS theory. In *Proc. of UAI'92*, pages 182–190, 1992.
- [10] A.S. Rao. AgentSpeak(L): BDI agents speak out in a logical computable language. In *Agents Breaking Away*, pages 42–55. Springer, 1996.
- [11] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- [12] M. Wooldridge. *An Introduction to Multiagent Systems*. Wiley, New York, 2009.