Distributed Belief Revision vs. Belief Revision in a Multi-Agent environment: first results of a simulation experiment

Aldo Franco Dragoni, Paolo Giorgini, Marco Baffetti

Istituto di Informatica, Università di Ancona,
via Brecce Bianche, 60131, Ancona (Italy)
{dragon,giorgini}@inform.unian.it

Abstract. We propose a distributed architecture for belief revision-integration, where each element is conceived like a complex system able to exchange opinions with the others. Since nodes can be affected by some degree of incompetence, part of the information running through the network may be incorrect. Incorrect information may cause contradictions in the knowledge base of some nodes. To manage these contradictions, each node is equipped with a belief revision module which makes it able to discriminate among more or less credible information and more or less reliable information sources. Our aim is that of comparing on a simulation basis the performances and the characteristics of this distributed system vs. those of a centralised architecture. In this paper we report the first results of our experiments.

1 Introduction

Mason's distributed seismic recognition system [1] is an example of successful application of DAI techniques. Conceived during and after the age of the "Salt I" and "Salt II" treaties against the proliferation of nuclear weapons, the system main task is that of integrating information coming from a network of seismographs distributed over the globe in order to discriminate natural seismic events from underground nuclear explosions. Distributed Truth Maintenance System (DTMS) [2] is one of the theoretical backgrounds of the latest versions of the system. A limit of DTMS is that it presupposes the trustworthiness of the network's nodes. For cases in which the nodes may be mutually inconsistent, the authors proposed a Distributed *Assumption-Based* Truth Maintenance System (DATMS) [3]. By supporting multiple contexts, DATMS was a step toward what they called "liberal belief revision policy": "it is better let agents stand by their beliefs based on their own view of the evidence".

We agree that ATMS [4] is central to belief revision. Trying to develope a method to perform belief revision in a multi-source enviroment (MSBR) [5], we realized that three relevant items are:

- maximal consistency of the revised knowledge base
- credibility of the information items
- reliability of the information sources [6].

Achieving maximal consistency is a symbolic process and can be accomplished by an ATMS. On the other hand, the credibility of the beliefs and the reliability of the agents can hardly be estimated without numerical processing (albeit we recognize the importance of qualitative methods whenever numbers make no sense).

MSBR (section 2) can be regarded as a way to integrate data coming from (eventually conflicting) sensors, databases, generic knowledge repositories and even witnesses during a trial or an inquiry [7].

However, MSBR is not a distributed approach since the data integration is attained in a centralized way. In a previous paper [8] we introduced the idea of Distributed Belief Revision (DBR). With DBR, nodes interact with each others in order to accomplish their own task. Since they can be affected by some degree of incompetence (eventually also insincerity), part of the information running through the network may be incorrect. Occasionally, incorrect information may cause contradictions in the knowledge base of some nodes. To manage these contradictions, *each* node is equipped with a MSBR module which makes it able to discriminate among more or less *credible* information items and more or less *reliable* information sources. DBR may be seen as a generalization of MSBR since it adopts the same mechanism but in a decentralized fashion. As we'll see in section 3, a major difference is that, being MSBR part of the node, also its rules and its numerical methods will be exposed to the other nodes' judgement.

Comparing features and performances of DBR w.r.t. MSBR can be done only on a simulation basis and on a common ground. Our testbed case is the task of extracting as much correct data as possible from a not updated and very unreliable database. This task can be accomplished under both, the MSBR and the DBR paradigms. Section 4 illustrates this point, presenting the first partial results of a simulation experiment that we are currently carrying on with the DBR architecture. We will conclude in section 5 by outlining what will be our work in the next few months.

2 A model for Belief Revision in a Multi-Source Environment (MSBR)

Derived from researches in multi-agent [9] and investigative domains [10], MSBR is a novel assembly of known techniques to the treatment of consistency and uncertainty. Let us recapitulate here the main ideas.

Defined as a symbolic model-theoretical problem [11-13,17], belief revision has also been approached both as a qualitative syntactic process [14,15] and as a numerical mathematical issue [16]. Trying to give a synoptic (although approximate) perspective of this composite subject, we begin by saying that both the cognitive state and the incoming information can be represented either as sets of weighted sentences or as sets of weighted possible worlds (the models of the sets of sentences). Weights can be either reals (normally between 0 and 1), representing explicitly the credibility of the sentences/models, or ordinals, representing implicitly the believability of the sentences/models w.r.t. the other ones. Essentially, belief revision consists in the redefinition of these weights in the light of the incoming information.

According to us, in a multi-agent environment, where information come from a variety of sources with different degrees of reliability, belief revision has to depart considerably from its original framework. Particularly, the principle of "priority of the incoming information" should be abandoned. While it is acceptable when updating the representation of an *evolving* world, that principle is not generally justified when revising the representation of a *static* situation. In this case, the chronological sequence of the informative acts has nothing to do with their credibility or importance. Another point is that changes should not be irrevocable. To make practical and useful belief revision in a multi-agent environment, we substitute the priority of the incoming information with the following principle[18].

Recoverability: any previously believed information item must belong to the current cognitive state if it is consistent with it.

We will achieve recoverability by imposing the maximal consistency of the revised cognitive state.

Along the paper we will represent *beliefs* as sentences of a propositional language *L*, with the standard connectives $\wedge$, $\vee$, $\rightarrow$ and $\neg$. Ξ is its set of propositional letters. Beliefs introduced directly by the sources are called *assumptions*. Those deductively derived from the assumptions are called *consequences*. Each belief is embodied in an ATMS node:

<*Identifier*, *Belief*, *Source*, *Origin Set* >

If the node represents a consequence, then *Source* (*S*) contains only the tag "`derived`" and *Origin Set* (*OS*) (we borrowed the name from [19]) contains the identifiers of the assumptions from which it has been derived (and upon which it ultimately depends). If the node represents an assumption, then *S* contains its source and *OS* contains the identifier of the node itself. We call *Knowledge Base* (*KB*) the set of the assumptions introduced from the various sources, and *Knowledge Space* (*KS*) the set of all the beliefs (assumptions + consequences). *KB* and *KS* grow monotonically since none of their nodes is ever erased from memory. Normally both contain *contradictions*. A contradiction is a pair of nodes as follows:

{<_, α, _, _>, <_, $\neg\alpha$, _, _>}

Since propositional languages are decidable, we can find all the contradictions in a finite amount of time. Inspired by Kleer, we define "nogood" a minimal inconsistent subset of *KB*, i.e., a subset of *KB* that supports a contradiction or an incompatibility and is not a superset of any other nogood. Dually, we define "good" a maximal consistent subset of *KB*, i.e., a subset of *KB* that is neither a superset of any nogood nor a subset of any other good. Each good has a corresponding "context", which is the subset of *KS* made of all the nodes whose *OS* is a subset of that good. Any node belongs to more than one context. Managing multiple contexts makes it possible to compare the credibility of different goods as a whole rather than confronting the credibility of single beliefs.

Procedurally, our method of belief revision consists of four steps:

S1. Generating the set *NG* of all the nogoods and the set *G* of all the goods in *KB*

S2. Defining a credibility ordering $\leq_{KB}$ over the assumptions in *KB*

S3. Extending $\leq_{KB}$ into a credibility ordering $\leq_G$ over the goods in G

S4. Selecting the preferred good CG with its corresponding context CC.

S1 deals with consistency and works with the symbolic part of the beliefs. Given an inconsistent *KB*, *G* and *NG* are duals: if we remove from *KB* exactly one element for each nogood in *NG*, what remains is a good. Let us recall here the definition of "minimal hitting-set". If *F* is a collection of sets, a *hitting-set* for *F* is a set $H \subset \bigcup_{S \in F} S$ such that $H \cap S \neq \emptyset$ for each $S \in F$. A hitting-set is *minimal* if none of its proper subsets is a hitting-set for *F*. It should be clear that *G* can be found by calculating all the minimal hitting-sets for *NG*, and keeping the complement of each of them w.r.t. *KB*. We adopt the set-covering algorithm described in [20] to find *NG* and the corresponding *G*, i.e., to perform S1.

S2 deals with uncertainty and works with the "weight", or the "strenght", of the beliefs. A question is: should $\leq_{KB}$ be a symbolic (qualitative, implicit) ordering (relative classification without numerical weights) or should it be a numerical (quantitative, explicit) one?. The first approach seems closer to the human cognitive behavior (which normally refrains from numerical calculus). The second approach seems more informative because it takes into account not just relative positions but also the gaps between the degrees of credibility of the various information items. In our MSBR we adopted the Demster-Shafer Theory of Evidence, in the special guise in which Shafer and Srivastava apply it to the "auditing" domain [22] since it provides a very intuitive tool to combine evidences from different sources, the Dempster Rule of Combination. The main reference for the Theory of Evidence is still [23]; a general discussion can be found in [24]. We recapitulate here the main concepts, definitions and rules, as they have been exploited in our MSBR.

To begin with, we introduce two data structures: the *reliability set* and the *information set*. Let $S=\{S_1,\ldots,S_n\}$ be the set of the sources and $I=\{I_1,\ldots,I_m\}$ be the set of the information items given by these sources. Then:

- *reliability set* = $\{<S_1, R_1>,\ldots,<S_n, R_n>\}$, where R_i (a real in [0,1]) is the reliability of S_i, interpreted as the "a priori" probability that S_i is reliable.
- *information set* = $\{<I_1, Bel_1>,\ldots,<I_m, Bel_m>\}$, where Bel_i (a real in [0,1]) is the credibility of I_i.

The reliability set is one of the two inputs of the belief-function formalism (see figure 3). The other one is the set $\{<S_1, s_1>,\ldots,<S_n, s_n>\}$, where s_i is the subset of I made of all the information items given by S_i. The information set is the main output of the belief-function formalism. Figure 2 presents the I/O mechanism applied to the case in the previous example.

Let us see now how the mechanism works. Remember that Ξ denotes the set of the atomic propositions of *L*. The power set of Ξ, $\Omega=2^{\Xi}$, is called *frame of discernment*. Each element ω of Ω is a "possible world"or an "interpretation" for *L* (the one in which all the propositional letters in ω are true and the others are false). Given a set of sentences $s \subseteq I$ (i.e., a conjunction of sentences), $[s]$ denotes the interpretations which are a model for all the sentences in *s*.

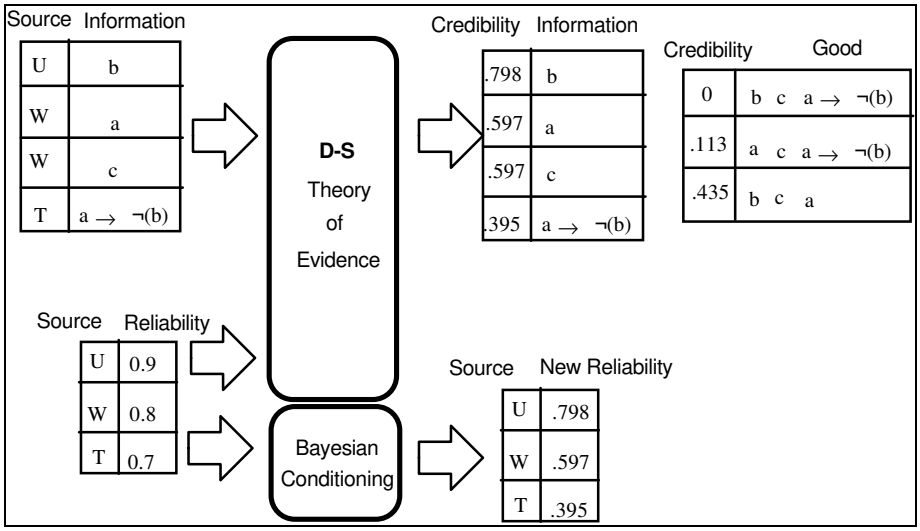

Fig. 2. Basic I/O of the belief-function mechanism

The key assumption with this multi-source version of the belief function framework is that *a reliable source cannot give false information, while an unreliable source can give correct information*; the hypothesis that "S_i is reliable" is compatible only with $[s_i]$, while the hypothesis that "S_i is *un*reliable" is compatible with the entire set Ω. Each S_i gives an *evidence* for Ω and generates the following *basic probability assignment* (*bpa*) m_i over the elements X of 2^{Ω}:

$$m_i(X) = \begin{cases} R_i & \text{if } X = [s_i] \\ 1 - R_i & \text{if } X = \Omega \\ 0 & \text{otherwise} \end{cases}$$

All these *bpa*s will be then combined through the Dempster Rule of Combination:

$$m(X)=m_1(X)\otimes\ldots\otimes m_n(X)= \frac{\sum_{X_1\cap\ldots\cap X_n=X} m_1(X_1)\cdot \ \ldots \ \cdot m_n(X_n)}{\sum_{X_1\cap\ldots\cap X_n\neq\varnothing} m_1(X_1)\cdot \ \ldots \ \cdot m_n(X_n)}$$

From the combined *bpa m*, the credibility of a set of sentences *s* is given by:

$$Bel(s) = \sum_{X\subseteq[s]} m(X)$$

In figure 2, we see another output of the mechanism, obtained through Bayesian Conditioning: the set $\{<S_1, NR_1>,\ldots,<S_n, NR_n>\}$, where NR_i is the new reliability of S_i. Following Shafer and Srivastava, we defined the "a priori" reliability of a source as the *probability* that the source is reliable. These degrees of probability are "translated" into belief-function values on the given pieces of information. However, we may also want to estimate the sources' "a posteriori" degree of reliability from the cross-examination of their evidences. To be congruent with the "a priori"

reliability, also the "a posteriori" reliability must be a *probability* value, not a belief-function one. This is the reason why we adopt the Bayesian Conditioning instead of the Theory of Evidence to calculate it. Let us see in detail how it works here. Let us consider the hypothesis that only the sources belonging to $\Phi \subseteq S$ are reliable. If the sources are independent, then the probability of this hypothesis is $R(\Phi)=\prod_{S_i \in \Phi} R_i \cdot \prod_{S_i \notin \Phi} \left(1 - R_i\right)$. We could calculate this "combined reliability" for any subset of S. It holds that $\sum_{\Phi \in 2^S} R(\Phi)=1$. Possibly, the sources belonging to a certain Φ cannot all be considered reliable because they gave contradictory information, i.e., a set of information items s such that $[s]=\emptyset$. In this case, the combined reliabilities of the remaining subsets of S are subjected to the Bayesian Conditioning so that they sum up again to "1"; i.e., we divide each of them by "1- $R(\Phi)$". In the case where there are more subsets of S, say $\Phi_1, \ldots, \Phi_l$, containing sources which cannot all be considered reliable, then $R(\Phi)=R(\Phi_1)+ \ldots +R(\Phi_l)$. We define the revised reliability NR_i of a source S_i as the sum of the conditioned combined reliabilities of the "surviving" subsets of S containing S_i. An important feature of this way to recalculate the sources' reliability is that if S_i is involved in contradictions, then $NR_i \leq R_i$, otherwise $NR_i=R_i$.

The main problem with the belief funcion formalism is the computational complexity of the Dempster's Rule of Combination; the straight-forward application of the rule is exponential in the frame of discernment (number of propositional letters of L, that is smaller than the number of information items in KB) and the number of evidences. However, much effort has been spent in reducing the complexity of the Dempster's Rule of Combination. Such methods range from "efficient implementations" [25] to "qualitative approaches" [26] through "approximate techniques" [27].

S3 also deals with uncertainty, but at the goods' level, i.e. it extends the ordering $\leq_{KB}$, defined at S2 on the assumptions, into an ordering $\leq_G$ onto the goods. This method could take into account either the ordinal weights (positions) of the assumption in $\leq_{KB}$, or their numerical weights. Let G' and G" be two goods of G. Among the former methods, Benferhat et al. [15] suggest the following three ways:

- *best-out method*. Let g' and g" be the most credible assumptions (according to $\leq_{KB}$), respectively, in KB\G' and KB\G". Then G"$\leq_G$G' iff $g' \leq_{KB} g$".
- *inclusion-based method*. Let G'_i and G''_i denote the subsets of, respectively, G' and G", made of all the assumptions with a priority i in $\leq_{KB}$. G"$\leq_G$G' iff a degree i exists such that $G'_i \supset G''_i$ and for any $j>i$, $G'_j = G''_j$. The goods with the highest priority obtained with this method are the same obtainable with the best-out method.
- *lexicographic method*. G"$<_G$G' iff a degree i exists such that $|G'_i|>|G''_i|$ and for any $j>i$ $|G'_j|=|G''_j|$, and G"$=_G$G' iff for any j, $|G'_j|=|G''_j|$.

Although the "best-out" method is easy to implement, it is also very rough since it discriminates the goods by confronting only two assumptions. The lexicographic method could be justified in some particular application domains (e.g. diagnosys).

The inclusion-based method yields a partial ordering $\leq_G$. In the special case in which $\leq_{KB}$ is complete and strict, the three methods produces the same ordering $\leq_G$. In this case, the algorithm in figure 3 (adapted from [5]) is a natural implementation of the method.

```
INPUT: Set of goods G
OUTPUT: G ordered by ≤G (when ≤KB is strict and complete)

G1 := G
repeat
  stack := KB ordered by ≤KB (most credible on the top)
  G2 := G1
    repeat
      pop an assumption A from stack
      if there exists a good in G2 containing A
        then delete from G2 the goods not containing A
    until G2 contains only one good G
  put G in reverse_ordered_goods
  delete G from G1
until G1 = ∅
return reverse(reverse_ordered_goods)
```

Fig. 3. An algorithm to sort the "goods" of the knowledge base

Among the "quantitative" explicit methods to perform S3, ordering the goods according to the average credibility of their elements seems reasonable and easy to calculate. A main difference w.r.t. the previous methods is that the preferred good(s) may no longer necessarily contain the most credible piece(s) of information.

The belief-function formalism is able to attach directly a degree of credibility to any good *g*, bypassing S3 in our framework. The problem is that if a good contains only part of the information items supplied by a source, then its credibility is null (see [9] for an explanation). Unfortunately, the event is all but infrequent, so that often the credibility of all the goods is null. This is the reason why we adopt the best-out or the "average" method to perform S3 in MSBR.

S4 consists of two substeps.

- Selecting a good *CG* from *G*. Normally, *CG* is the good with the highest priority in $\leq_G$. In case of ties, *CG* might be either one of those with the same highest priority (randomly selected) or their intersection (see [15] and [21]). This latter case means rejecting all the conflicting but equally credible information items. The result is not a good (it is not maximally consistent) and thus implies rejecting more assumptions than necessary to restore consistency. We believe that this could be avoided by simply considering $\leq_G$ as a *primary* ordering that could be combined with whatsoever user-defined classification to reduce or eliminate the cases of ties.

- Selecting from *KS* the derived sentences whose *OS* is subset of *CG* to build *CC*. We could relax the definition of *OS* to that of "the set of assumptions used (but not all necessary) to derive the consequence". This is easier to compute and does not have pernicious repercussions; the worst that can happen is that, this relaxed *OS* being a superset of the real one, is not necessarily a subset of *CG* (whenever the real *OS* is), and thus the consequence node is erroneously removed from *CC*.

3 A model for Distributed Belief Revision (DBR)

MSBR is not a distributed model since the integration/revision of the information is accomplished in a centralized way by an external module. Its simple behavior is described in figure 4.

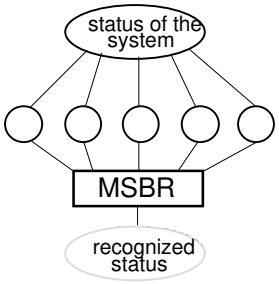

Fig. 4. Centralized Belief Revision (MSBR)

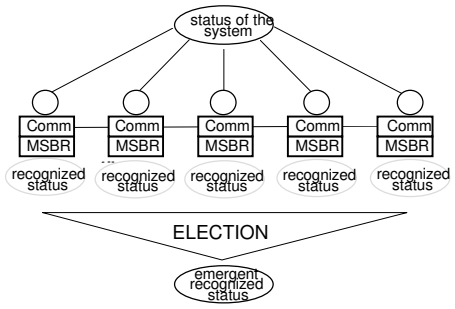

Fig. 5. Distributed Belief Revision (DBR)

Figure 5 shows what we mean by Distributed Belief Revision. Under normal operating conditions, we do not expect that the global output emergent from DBR will be better than the output of MSBR, neither for the quality nor for the quantity of the information provided. What we expect is that the distributed architecture will be more "robust" and "fault tolerant", i.e., it will be able to offer an acceptable output even in the case that some nodes are seriously compromised. On the other hand, nowadays DBR is a viable alternative to MSBR since the prices of hardware (CPU, RAM and, expecially, mass storage) and the communication costs have fallen dramatically.

In DBR nodes exchange information with each others. Thus we need to equip them with (at least) two modules:

1. a communication module (Comm), which deals with the three fundamental communication policies:
   a) choice of the *recipient* of the communication
   b) choice of the *argument* of the communication
   c) choice of the *time* of the communication
2. a model for belief integration/revision in a multi-source environment (MSBR)

In order to increase the realism of this distributed architecture, we introduced two fundamental assumptions:

1. nodes do not communicate to the others the sources from where they received the data, but they present themselves as completely responsible for the knowledge they are passing onto the others; a receiver considers the sender as the source of the information it is sending
2. nodes do not exchange opinions regarding the reliability of the other nodes with whom they got in touch.

With 1 we extend the scope of responsibility: a node is responsible not only for the information that it provides to the network as its original source, but also for the information that it receives from other nodes and, *retaining it credible*, passes it on to the others. With 2 we limit the range of useful information: an agent's opinion regarding the others' (and its own) reliability is drawn out from pure data regarding the knowledge domain under consideration, not from indirect opinions.

Among the various thinkable criteria to select the recipient of the communication, we see that two of them are worth noticing. Being guided by an "esprit de corps", one should offer its best information to the node it retain the least reliable one, with the aim of increasing its reliability. But the same collaborative spirit could lead to the opposite conclusion: one should send its best information item to the most reliable node since, if it will be recognized also by the others as the most reliable one, then that information item will be spreaded over all the group. The latter criterion implies that unreliable nodes will gradually be isolated from the rest of the group.

By comparing its opinion with the others' ones, each node produces its own reconstruction of the status of the system. The effects of the others' opinions depend on the rules adopted by the MSBR module.

Although not necessary, we may want to extract from the network an emergent global reconstruction of the system's status. To preserve the decentralized nature of DBR, the global opinion shoud be synthetized not by an external supervisor/decisor, but through some form of *election*: the group elects what it believes the global output to be returned to the external world should be. However, nothing prevents a user to get his/her information directly from a single node's output, since the election does not change the node's personal opinions. If the various cognitive states are quite similar, then the global output cannot differ very much from each node's one. According to us, the similarity between the opinions of the various nodes could be taken as a parameter to evaluate the quality of the Comm and MSBR modules.

The election of the group's emergent output could be done in several ways. We have no room here to explore sufficiently this matter, however at the extreme positions we see two distinct kinds of election:

1. "data driven" election: the candidates are information items; only the winners will be part of the global output (direct synthesis of the global output)
2. "node driven" election: the candidates are nodes of the network; only the winners will be charged to make up the global opinion (synthesis "by proxy" of the global output)

Many strategies can be conceived by mixing these two kinds of election.

We believed that the comparison of the characteristics and performances of the two paradigms illustrated in Fig. 4 and Fig 5 could be done only on a simulation basis.

4 The simulation experiment

The task given to the group is that extracting as much correct knowledge as possible from a corrupted knowledge repository. The knowledge repository is trivially implemented as a couple of databases containing the same quantity of information items: a database holds correct information, the other one contains the negations of the information items in the correct database. Nodes cannot distinguish the two databases. Each node is characterized by a degree of "capacity" (between 0 and 1) that will be adopted as the frequency with which it accesses (unconsciously) the correct database. Fig 6 should clarify the structure of the experiment.

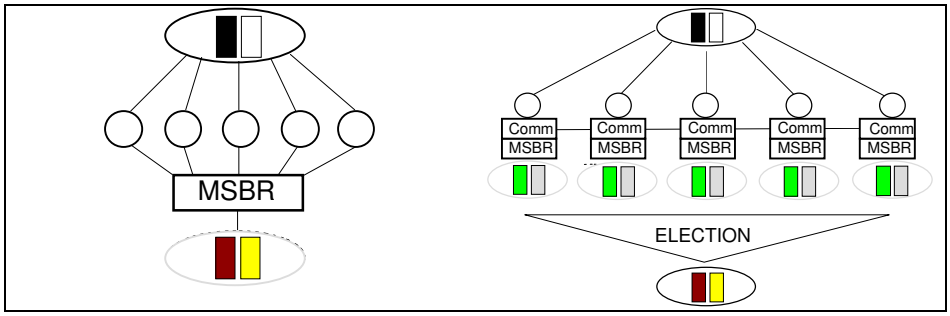

Fig. 6. Structure of the experiment: comparison between MSBR and DBR

Our final goal is that of comparing the ability of the centralized and of the distributed architectures in reconstructing the two databases. However, in this initial phase of our study, we have focused on the distributed approach because it revealed to be much more complex and tangled than expected.

Unfortunately, due to the length of a simulation (in the average case 12 hours on a Sparkstation 2 with 32 MB of RAM) we are not still able to report the effects of the various strategies (we have conceived) to perform the "election" phase (reconstruction of the global output from the opinions of the various nodes). For the moment we have limited ourselves to study the effects that interaction had on each node's cognitive state; i.e., how much the cognitive state differs from the case that the node would have not interacted with the others. This study should return a measure of (eventual) convenience for each node of having been part of the network.

Nodes access the databases once per simulation cycle.

Regarding the communication policies ("peer-to-peer" communication), the results presented in this paper refer to the most aleatory case in which:

a) the *recipient* of the communication is selected randomly
b) the *argument* of the communication is a randomly selected datum in the preferred good

c) each node communicates once per simulation cycle

In each node' cognitive state KB, the preferred good G is taken as the reconstruction of the correct database, while KB/G is taken as the reconstruction of the corrupted database.

For each node we evaluated three parameters: its average *reliability*, the *quality* and the *quantity* of the data in its preferred good. Quality and quantity are evaluated as differences w.r.t. the case without interaction.

Quality = $Q - Q_{\text{without communication}}$

where $Q = \dfrac{|\text{true propositions in G}| + |\text{false propositions outside G}|}{|\text{propositions in KB}|}$

Quantity = $|\text{true propositions in G}| - |\text{true propositions in G}|_{\text{without communication}}$

Reliability of the i^{th} node $R_i = \dfrac{\sum_{k=1}^{|\text{nodes}|} r_{ki}}{|\text{nodes}|}$ where r_{ki} is the reliability of the i^{th} node, estimated by the k^{th} one.

Due to severe complexity and duration problems, we fixed at five the number of nodes of the network. Each of the two databases contains 100 propositions. From some tentatives with 40 accesses for each node, we realized that 25 accesses were sufficient to reach stable values. To reduce the effects of casuality, each of the sixteen simulations has been repeated 20 times. We calculated the average of each parameter, for each cycle, over the twenty repetitions.

Experiment A

We began by making two series of simulations. In the first one we gave to the five nodes the same capacity. We made six simulations, from capacity 0 to capacity 1 (fig.7). In the second series (ten simulations) we gave to the nodes the capacities reported in fig.8.

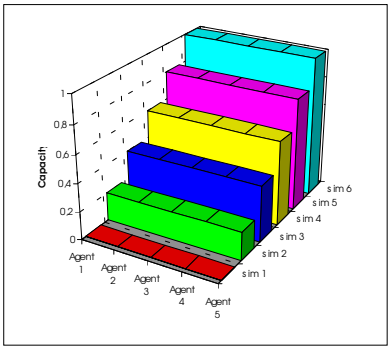

Fig. 7

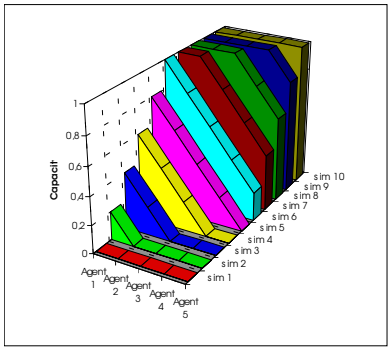

Fig. 8

The results obtained can be summarised as follows.

Quality

Interaction increases the quality of an incapable node' cognitive state and decreases the quality of a capable one. A typical trend is reported in fig.9, which refers to the sim 5 in fig.8 (decreasing capacity from 1 to 5). The average quality (Av) stands at zero. This means that if there are few incapable nodes, then they gain much in correctness while the others lose very little.

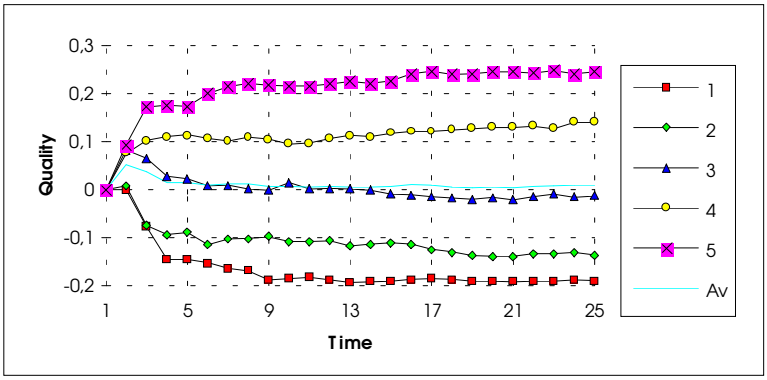

Fig.9

Quantity

Interaction always increases quantity. However, incapable nodes gain more than the capable ones. A typical trend is reported in fig.10 (sim 5 fig. 8; decreasing capacity from 1 to 5).

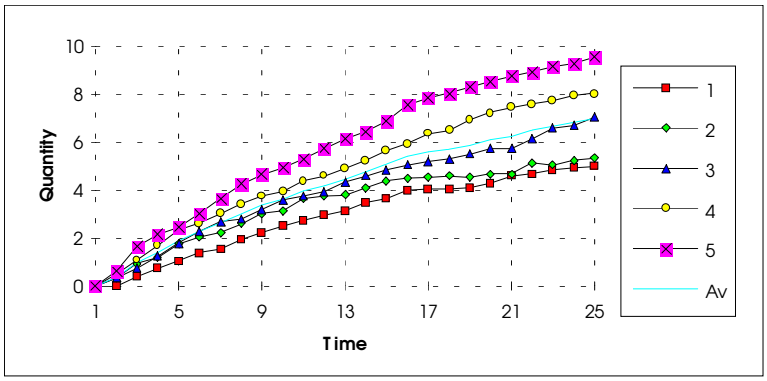

Fig.10

Reliability

Reliabilities follow typical trends as the one reported in fig.11 (sim 7 in fig.8; decreasing capacity from 1 to 5). We see that all the nodes lose reliability, but there is what we called "majority effect". If the average capacity of the group is greater than 0.5, then the capable nodes lose less than the incapable ones.

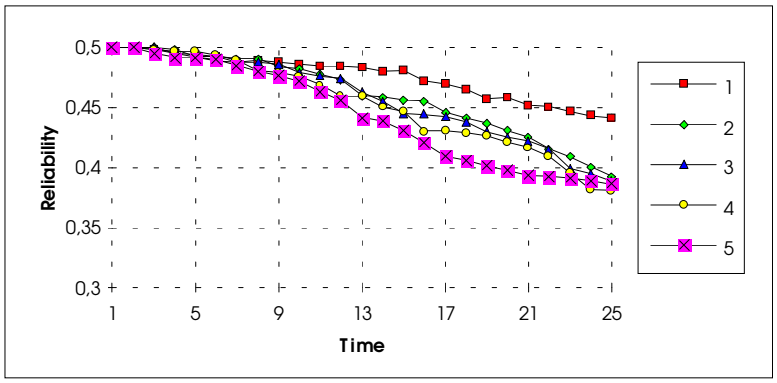

Fig.11

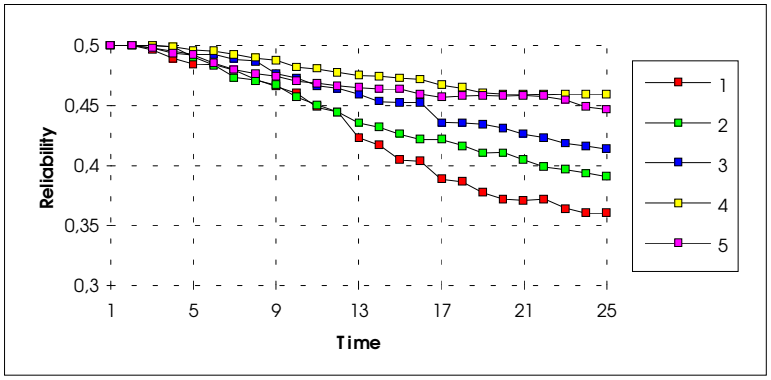

Fig.12

On the opposite, if the average capacity is less than 0.5, than the capable nodes lose more than the incapable ones, as in fig.12 (sim 3 fig.8). In this case, since the capable agents are in minority, they are regarded as unreliable by the others. The result is that the estimated reliability is quite inversely proportional to the real capacity.

These results were expected, since this group looks like a class without teacher, or a scientific community without experimental evidence. They are positive in the sense that the members of the group acquire knowledge without loosing (averagely) in correctness.

Experiment B

Subsequently, we introduced a "teacher" node with the following features:

- capacity = 1 (it accesses only the correct database)
- for each k, $r_k = 1$ (all the nodes know that the teacher is absolutely reliable)
- it transmits but doesn't receive information (the teacher's cognitive state will not be contaminated by the others)

We repeated the simulations with the teacher and five nodes as in the experiment A. Results may be summarised as follows.

Quality

The main difference w.r.t. the case without teacher is that now the average quality is positive and lightly increasing. This is due to the presence of the teacher which, occasionally (communications took place randomly) gave to the others the possibility to choose in the right way. A typical trend is showed in fig.13 (capacity decreases from node 1 to node 5).

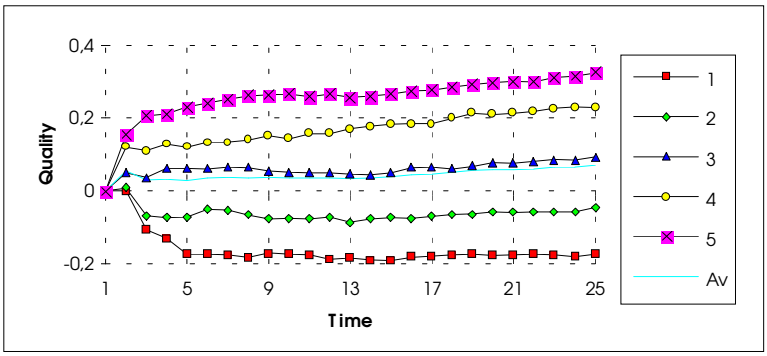

Fig.13

Quantity

Even in this case, interaction increases the amount of correct data in the cognitive state of every node, and the gain is inversely proportional to the capacity of the node. The effect of the teacher is that the gain is higher than before for any node. A typical trend is showed in fig. 14 (decreasing capacity from 1 to 5).

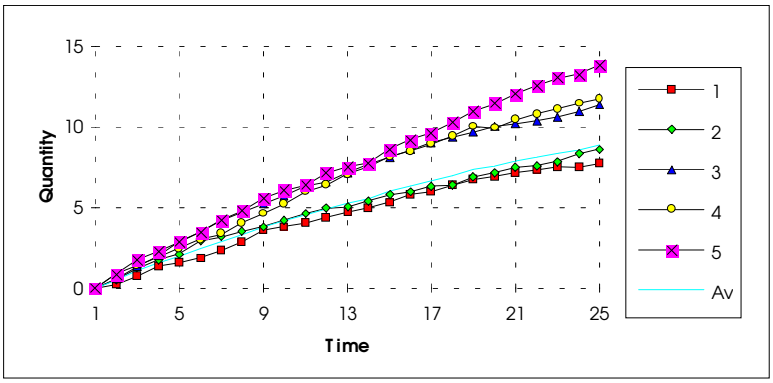

Fig.14

Reliability

The correspondence between estimated reliability and real capacity holds when the average capacity is greater than 0.5. A typical trend is showed in fig. 15 (decreasing capacity from 1 to 5). Under that value the situation becomes confused. What we

called “majority effect” in the experiment A is no longer appreciable when there is a teacher in the group.

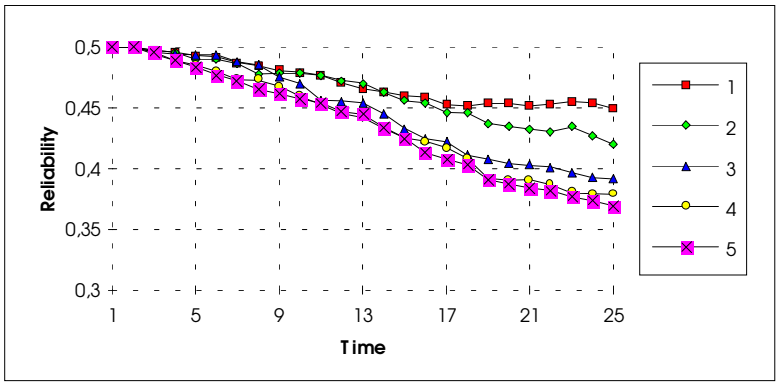

Fig.15

Experiment C

In the previous experiments, the characteristics of the nodes were static; their capacity did not change during the simulation. We were curious to see if the group would have been able to realize that some of its member changed their degree of capacity, and how long would have it taken to the group to be aware of the change. After several simulations with a node decreasing or increasing its capacity, we realized that only high quality groups were able to perceive the change. For groups with an average capacity less than 0.6, the situation at the time of the change was sufficiently cahotic to hide it. This implies that only decreases of capacity will be perceived by the group. The graph in fig. 16 reports the “quality” trends of five nodes with capacity 0.98, where the fifth one, at the 15th simulation cycle, decreased its capacity to 0.8.

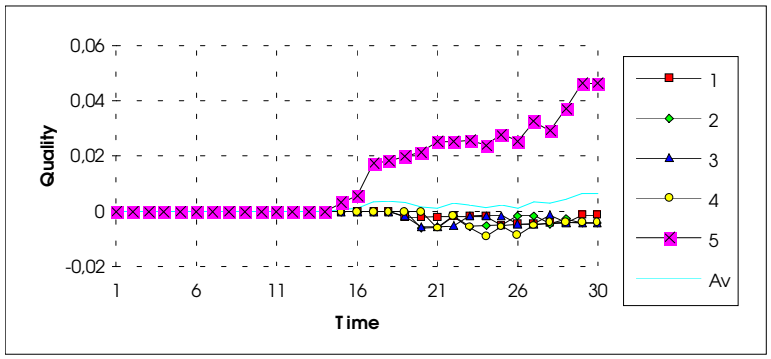

Fig.16

The graph in fig. 17 reports the “quantity” trends of five nodes with capacity 0.8 where the fifth one, at the 15th simulation cycle, decreased its capacity to 0.6.

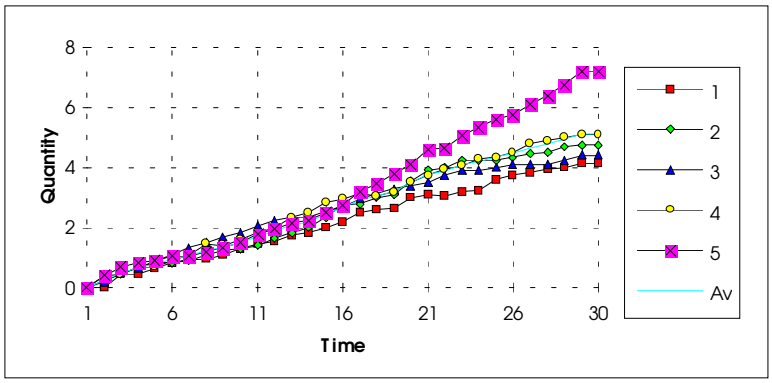

Fig.17

The graph in fig. 18 reports the "reliability" trends of five nodes with capacity 0.98 where the fifth one, at the 15th simulation cycle, decreased its capacity to 0.8.

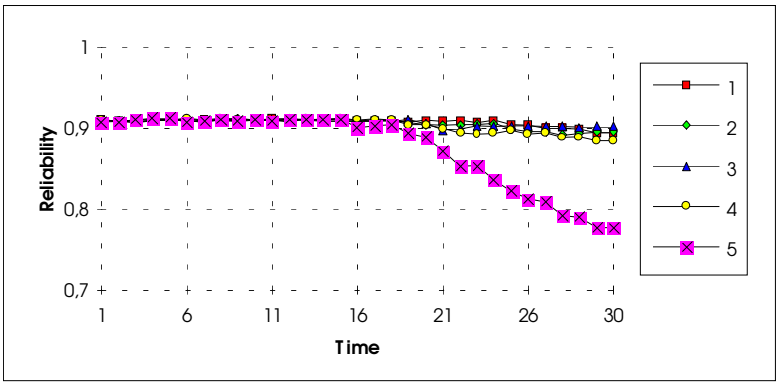

Fig.18

As an example of a low quality group, fig. 19 reports the "reliability" trends of five nodes with capacity 0.6 where the fifth one, at the 15th simulation cycle, decreased its capacity to 0.4. It is evident that the group was not able to perceive the change.

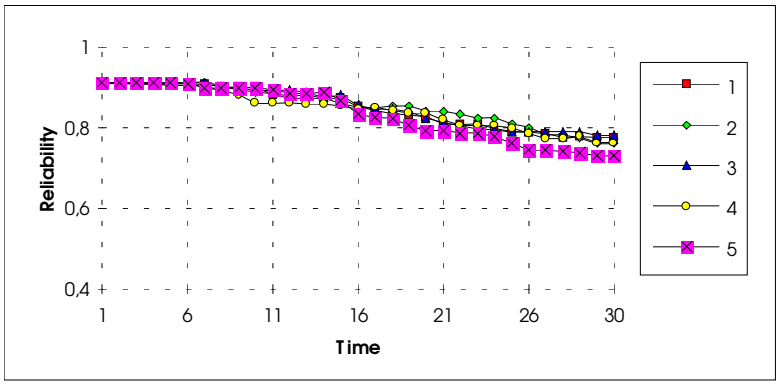

Fig.19

We also tried to study the effects of the deepness of the reciprocal acquaintance: the longer the period spent togheter, the more reluctant should be the nodes in changing their opinions regarding their companion who decreased its capacity. As a matter of fact, fig. 20 shows the "reliability" trend of the node who changed its capacity from 0.98 to 0.8 at different time points: t=0, 5, 10, 15, 20 and 25. We can see that the later the change, the lighter the curve's slope: the group which appreciated the node's reliability for a longer period has more inertia to change its opinion regarding that node.

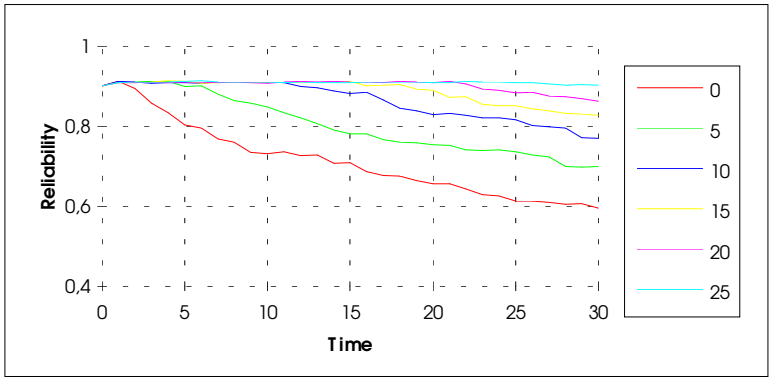

Fig.20

Conclusions

While we are writing, we are very far from the conclusion of our project. We have still to try the effects of specific communication strategies on the formation of the various nodes' cognitive state. Furthermore, and expecially, we have still to make experiments with the different kinds of election which should form the emergent output of the group from the member's different cognitive states. Finally, we have still to make experiments in solving the same task with the centralized architecture (MSBR) to compare its output with the one emergent by the distributed architecture.

The reason why we are so late is not only the duration of the simulations (each of the three experiments described in this paper took almost a mounth of our time), but also the fact that ideas have not been always so clear. The distributed architecture revealed to be much more tangled than expected. For instance, when we started, we were not interested in the dynamic cognitive behaviour of the group; we realized its importance only while we were working. We left out other intricate questions regarding the structure of the group and the nature of authority (although some of these questions will be treaten when we'll study the election phase). We could compare the performances of the two architectures under a MSBR model based on different approaches to the treatment of uncertainty (probabilistic, possibilistic, ...), and, indeed, may be that we'll do it in the future.

For the moment, the conclusion is that the MSBR module we adopted here shows interesting properties even under a multi-agent perspective. The global cognitive

behaviour of the group seems sociologically intuitive and foreseeable (hence, acceptable).

References

[1] Mason, C., An Intelligent Assistant for Nuclear Test Ban Treaty Verification, IEEE Expert, vol 10, no 6, 1995.

[2] Huhns, M. N., Bridgeland, D. M.: Distributed Truth Maintenance. In Dean, S. M., editor, Cooperating Knowledge Based Systems, pages 133-147. Springer-Verlag, 1990.

[3] Cindy L. Mason and Rowland R. Johnson, DATMS: A Framework for Distributed Assumption Based Reasoning, in L. Gasser and M. N. Huhns eds., Distributed Artificial Intelligence 2, Pitman/Morgan Kaufmann, London, pp 293-318, 1989.

[4] de Kleer J., An Assumption Based Truth Maintenance System, in Artificial Intelligence, 28, pp. 127-162, 1986.

[5] Dragoni A.F., A Model for Belief Revision in a Multi-Agent Environment, in Werner E. and Demazeau Y. (eds.), Decentralized A. I. 3, North Holland Elsevier Science Publisher, 1992.

[6] Dragoni A.F., Belief Revision: from theory to practice, to appear on "The Knowledge Engineering Review", Cambridge University Press, 1997.

[7] Dragoni A.F., Ceresi, C. and Pasquali, V., A System to Support Complex Inquiries, in Proc. of the "V Congreso Iberoamericano de Derecho e Informatica", La Habana, 6-11 march 1996.

[8] A.F. Dragoni, P. Giorgini and P. Puliti, Distributed Belief Revision vs. Distributed Truth Maintenance, in Proc. 6th IEEE Conf. on Tools with A.I., IEEE Computer Press, 1994.

[9] A.F. Dragoni, P. Giorgini, "Belief Revision through the Belief Function Formalism in a Multi-Agent Environment", Third International Workshop on Agent Theories, Architectures and Languages, LNAI n°, Springer-Verlag, 1997.

[10] Dragoni, A.F., Maximal Consistency, Theory of Evidence and Bayesian Conditioning in the Investigative Domain, to appear on the "International Journal on Artificial Intelligence and Law", 1997.

[11] Alchourrón C.E., Gärdenfors P., and Makinson D., *On the Logic of Theory Change: Partial meet Contraction and Revision Functions*, in *The Journal of Simbolic Logic*, 50, pp. 510-530, 1985.

[12] P. Gärdenfors, Knowledge in Flux: Modeling the Dynamics of Epistemic States, Cambridge, Mass., MIT Press, 1988.

[13] P. Gärdenfors, Belief Revision, Cambridge University Press, 1992.

[14] W. Nebel, Base Revision Operations and Schemes: Semantics, Representation, and Complexity, in Cohn A.G. (eds.), Proc. of the 11th European Conference on Artificial Intelligence, John Wiley & Sons, 1994.

[15] Benferhat S., Cayrol C., Dubois D., Lang J. and Prade H., Inconsistency Management and Prioritized Syntax-Based Entailment, in Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence, pp. 640-645, 1993.

[16] Dubois D. and Prade H., A Survey of Belief Revision and Update Rules in Various Uncertainty Models, in International Journal of Intelligent Systems, 9, pp. 61-100, 1994.

[17] Williams M.A., *Iterated Theory Base Change: A Computational Model*, in *Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence*, pp. 1541-1547, 1995.

[18] Dragoni A.F., Mascaretti F. and Puliti P., *A Generalized Approach to Consistency-Based Belief Revision*, in Gori, M. and Soda, G. (Eds.), *Topics in Artificial Intelligence*, LNAI 992, Springer Verlag, 1995.

[19] Martins J.P. and Shapiro S.C. (1988), *A Model for Belief Revision*, in «Artificial Intelligence», 35, pp. 25-79.

[20] R. Reiter, A Theory of Diagnosis from First Principles, in Artificial Intelligence, 53, 1987.

[21] Benferhat S., Dubois D. and Prade H., *How to infer from inconsistent beliefs without revising?*, in *Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence*, pp. 1449-1455, 1995.

[22] Shafer G. and Srivastava R., The Bayesian and Belief-Function Formalisms a General Perpsective for Auditing, in G. Shafer and J. Pearl (eds.), Readings in Uncertain Reasoning, Morgan Kaufmann Publishers, 1990.

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

[24] Shafer G., *Belief Functions*, in G. Shafer and J. Pearl (eds.), *Readings in Uncertain Reasoning*, Morgan Kaufmann Publishers, 1990.

[25] Kennes, R., *Computational Aspects of the Möbius Transform of a Graph*, IEEE Transactions in Systems, Man and Cybernetics, 22, pp 201-223, 1992.

[26] Parson, S., *Some qualitative approaches to applying the Demster-Shafer theory*, Information and Decision Technologies, 19, pp 321-337, 1994.

[27] Moral, S. and Wilson, N., *Importance Sampling Monte-Carlo Algorithms for the Calculation of Dempster-Shafer Belief*, Proceeding of IPMU'96, Granada, 1996.