Paper presented at the RAMS'99
(Reliability and Maintainability Symposium)
Washington, January 1999

Assessment of a Safety-Critical System Including Software: A Bayesian Belief Network for Evidence Sources

Marc Bouissou • EDF, DER/ESF • Clamart
Fabrice Martin • EDF, DER/PEL • Clamart
Alain Ourghanlian • EDF, DER/CCC • Chatou

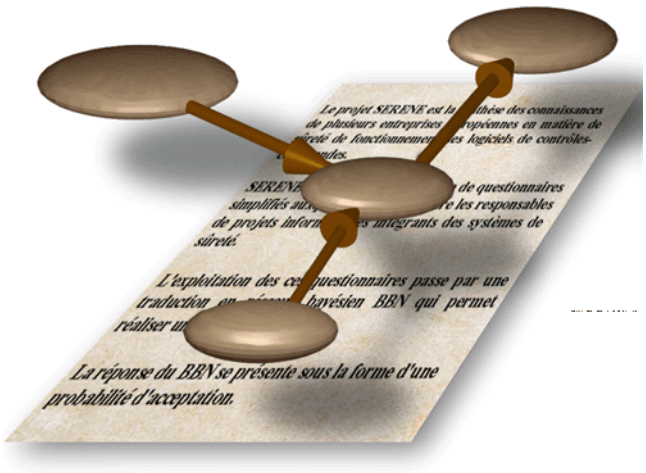

Assessment of a Safety-Critical System Including Software: A Bayesian Belief Network for Evidence Sources

✣ Marc Bouissou • EDF, DER/ESF • Clamart
Fabrice Martin • EDF, DER/PEL • Clamart
Alain Ourghanlian • EDF, DER/CCC • Chatou

Key Words: safety case, safety argument,Bayesian Belief Network, software reliability, SERENE, safety critical system.

SUMMARY & CONCLUSIONS

Assessment of safety critical systems including software cannot rely only on conventional techniques, based on statistics and dependability models. In such systems, the predominant faults usually are *design* faults, which are very hard to predict. Therefore, the assessment can only be qualitative, and is performed by experts, who take into account various evidence sources.

The aim of the SERENE European project is to improve the understandability, and repeatability of such assessments, thanks to a representation of the expert's reasoning by a mathematical model (a Bayesian Belief Network). The subject of this paper is the presentation of the BBN built by EDF to model one of its assessment approaches, valid for the products for which EDF writes the requirements specification, and then monitors the development made by an external supplier.

No doubt that, before it yields reliable forecasts, this kind of model will require many years of calibration, by comparison between the predictions it gives, and the real, observed safety level of the evaluated systems. However, we think that in the short term, they can bring a rationale in the discussions between experts. They will also help in determining which are the most influential variables in the design process of a system, which is a necessary prerequisite for setting up any kind of field experience collection.

1. INTRODUCTION

During the recent years, there has been a trend to replace conventional electro-mechanical systems for the control of industrial plants with computer based systems. Although this kind of technology clearly has many advantages, it is not yet widely used in safety critical applications, such as reactor protection systems of nuclear power plants.

The explanation for this reluctance is the complexity of the safety assessment, and the licensing of these systems, due in particular to the characteristics of software.

One cannot claim that the impact of design faults is negligible when compared to the effect of hardware faults (which are much easier to estimate): in fact, it is usually predominant ! This is unfortunate, because the design faults are by far the hardest to predict. This is due to the fact that a digital system behaviour is subject to discontinuities: the slightest change in its inputs, or in its design can result in a dramatically different (and potentially catastrophic)behaviour.

For safety critical systems, software reliability (and, more generally, the reliability related to design faults) cannot be estimated by means such as reliability growth models, because the number of observed failures is (fortunately) far too low to allow any statistical processing of this data.

Thus, the only way to license such a system is to build a so-called 'safety case', or 'safety argument', which is a collection of all kinds of evidence related to both the development process and the final product. The kind of information to be collected, and the way to do it, are now relatively well defined, by means of template documents, and check-lists such as those given in the SHIP (Ref. 1) and CASCADE (Ref2) ESPRIT projects.

But so far, no method has been proposed in order to combine the diverse sources of evidence into an overall safety assessment. This crucial work entirely relies on the assessor's expertise.

Electricité de France is currently involved in an ESPRIT project called SERENE (see partners list in the acknowledgements) with the objective of building a method and a tool that could contribute to this 'high level' part of safety arguments, in order to **improve the repeatability of the assessments and make them more understandable**.

SERENE relies on the use of Bayesian Belief Networks, a formalism well known for its ability to reason on uncertain facts, to model and formalize the assessor's expertise.

The object of this paper is the result of the work performed at EDF, which was to build a BBN that could help EDF's assessors to weight properly the various sources of evidence, in order to arrive at a final judgement.

The paper is organised as follows:

- Section 2 gives a definition of a BBN, and shows on a tutorial example how it can be exploited,
- Section 3 very briefly presents the SERENE method,
- Section 4 describes how EDF develops and assesses its instrumentation and control (I&C) systems,
- Section 5 describes the BBN we built,
- Section 6 gives some results of calculations performed with the BBN.

2. WHAT IS A BAYESIAN BELIEF NETWORK?

BBNs (sometimes called Belief Networks, Causal Probabilistic Networks, Causal Nets, Probabilistic Cause-Effect Models, and Probabilistic Influence Diagrams) are a quickly expanding technology, in many areas where decision aids are needed in a context of uncertain knowledge about the « real world », for example in the medical, military and financial domains (Ref.3).

The theoretical development of BBNs dates back to the 1970s, but at that time, the lack of efficient algorithms and of useful tools impaired the development of practical applications.

Nowadays, a number of tools (among which is HUGIN explorer, the tool used by the SERENE partners) allow on the one hand, efficient calculations on an existing BBN, and on the other hand, the automatic construction of a BBN from a (sufficiently large) database of experimental facts. Although the latter kind of tools should still be considered as ongoing research, they already give very promising results, leading way to an extremely wide range of applications.

A BBN is a relatively simple mathematical object; however, we will present it here in a rather informal way, in order not to puzzle the readers who are totally unfamiliar with this concept.

Any reader interested in a comprehensive, mathematical presentation of BBNs and the algorithms necessary to process them should refer to Ref. 4, whereas somebody who wants to know more about the applications, and how to build a BBN starting from human expertise might rather refer to Ref. 5. The latter book contains a demonstration version of the HUGIN explorer tool, which is also available on the web, at the address: www.hugin.dk.

2.1 A simple modeling problem involving uncertainty

The small example we are going to use now to introduce the BBN concepts is adapted from the SERENE Method document of the SERENE project.

Let us imagine that we wish to model the following **knowledge**: « Bill and John go to work using different transportation means. John drives his car, whereas Bill goes by train. Bill rarely misses his train, and the train is nearly always on schedule, *except on strike days*. However, a train strike does not *imply* that Bill will definitely be late (he might leave early and drive). A train strike can also cause John to be late because traffic is heavier in that case. But, John is often late anyway, because he often oversleeps, and therefore a train strike only increases the likelihood of his lateness by a small amount. In the event of a train strike, John is less likely to be late than Bill. »

Now, given that knowledge, how could we model the following **inferences,** which seem quite reasonable?

1- If we know Bill is late, we believe that there is a train strike, and we think that John is (slightly) more likely to be late,

2- Suppose we find out that John is late. This evidence increases our belief in both of the possible causes (namely a train strike and John oversleeping). But if we learn that Bill is late too, we will be tempted to infer that a train strike occurred, and was the cause of John's lateness, which decreases our belief that he overslept.

2.2 A BBN to represent knowledge

A BBN is a graph (consisting of nodes and arcs) together with an associated set of probability tables, called node probability tables (NPT), since they have a one to one relationship with the nodes of the graph.

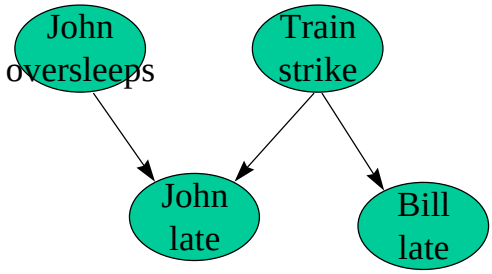

Figure 1 : A BBN representing the knowledge described in the text in § 2.1

The nodes represent discrete random variables. In our example, the four variables (see Figure 1 below) have only two possible values: 'True' and 'False'. The arcs represent causal relationships between variables. Since a train strike can cause Bill to be late we model this relationship by drawing an arc from the node 'Train strike' to the node 'Bill late'.

The key feature of BBNs is that they enable us to model and reason about uncertainty. In our example, in the BBN we model this by filling in a node probability table for each node. For the node 'Bill late' the NPT might look like this:

	'Train strike' = 'True'	'Train strike' = 'False'
Pr('Bill late' = 'True')	0.6	0.1

This is actually the **conditional probability of the variable 'Bill late' given the variable 'Train strike'** . The table gives only the probability of the event 'Bill late' = 'True' , since Pr('Bill late' = 'False') = 1- Pr('Bill late' = 'True') [1]. Informally, the particular values in this table tell us that Bill is very unlikely to be late normally, but if there is a train strike he is likely to be late (the probability is 0.6).

To reflect accurately our knowledge depicted above, we model the relationship between John's lateness and its two possible causes by the following probability table:

[1] Note that from that point, we will always give only the probability ofthe value 'True' for variables with only values 'True' and 'False'.

John oversleeps	True		False	
Train strike	True	False	True	False
Pr('John late' = 'True')	0.7	0.5	0.4	0.1

The probability tables associated with the nodes 'Train strike' and 'John oversleeps' are somewhat different in nature. These nodes have no 'parent' node in this model (we call them **root nodes**), and therefore we only have to assign a probability to each of their two possible values 'True' and 'False'. In fact, we will assume that Pr('Train Strike' = 'True') = 0.1, and that Pr('John oversleeps' = 'True') = 0.4.

There may be several ways of determining the probabilities of any of the tables. For example, we might be able to base the probability table of 'Train strike' on previously observed frequencies of days when there were train strikes. Alternatively, if no such statistical data is available we may have to rely on subjective probabilities entered by experts. The advantage of BBNs is that they allow to employ both subjective probabilities and probabilities based on statistical data in a unified framework.

2.3 Entering evidence in a BBN to make inferences

The **inferences** given in § 2.1 (and many more !) can be obtained automatically by ***calculations*** (according to probability theory) performed on the*BBN* we just described.

The BBN is in fact **a concise representation of the joint probability distribution of the BBN variables, which is the most complete information one can have about this set of random variables**. This joint probability distribution could, in this simple example, alternatively be given by a large unique table, giving the probabilities of the 2^4 possible combinations of values for all variables. Building this table would be an intractable task for a large BBN, but it is not necessary, since calculations can be performed directly on the BBN. For example, we might want to calculate the (unconditional) probability that Bill is late:

Pr(Bill late) = Pr(Bill late | train strike) * Pr(train strike) + Pr(Bill late | no train strike) * (1 - Pr(train strike))

= (0.6 * 0.1) + (0.1 * 0.9) = 0.15

This probability is an example of an "a priori" probability: it reflects our knowledge on the event 'Bill late' in the absence of information about the real state of any of the BBN nodes.

The following table recaps the a priori probabilities for all the nodes of our example BBN:

John oversleeps	Train strike	Bill late	John late
0.4	0.1	0.15	0.286

In fact, the most important use of BBNs is in revising probabilities in the light of actual observations of events.

- Suppose, for example, that we know there is a train strike. We can "enter this evidence" in the BBN, by setting Pr('Train strike' = 'True') to 1. Then the conditional probability tables already tell us the revised probability for Bill being late (0.6) and a calculation indicates that Pr(John being late) becomes 0.52, which is consistent with the sentence « In the event of a train strike, John is less likely to be late than Bill. » of § 2.1. If it was not, this would mean that the NPTs would have to be modified.

When we enter evidence and use it to update the probabilities in this way we call it **propagation**

- Suppose, now, that we do not know if there is a train strike but do know that Bill is late. Then we can enter the evidence that 'Bill late' = 'True' and we can use this observation to determine the (revised) probability that there is a train strike, and the (revised) probability that John will be late. In fact, the observation that Bill is late significantly increases the probability that there is a train strike (up from 0.1 to 0.4) and slightly increases the probability that John is late (from 0.286 to 0.364). This is the "formal" version for inference 1 in § 2.1.
- Here is a last propagation example, corresponding to inference 2 in § 2.1. Suppose we find out that John is late. This evidence increases our belief in both of the possible causes (namely a train strike and John oversleeping). Specifically, applying Bayes theorem yields a revised probability of train strike of 0.182 (up from the prior probability of 0.1) and a revised probability of John oversleeping of 0.727 (up from the prior probability of 0.4). Therefore, if we had to bet on it, our money would be firmly on John oversleeping as the more likely cause. Now suppose we also discover that Bill is late. Entering this evidence and applying Bayes yields a revised probability of 0.571 for a train strike and 0.637 for John oversleeping, which indicates, that in fact, the two causes are practically equally likely.

This very simple example shows that using a simple BBN model enables us to make inferences which are much more precise than what would have been possible by using just sentences with expressions such as « very likely », « unlikely », « slightly increases »... and so on. *The same remark is true when we model an expert's reasoning in a safety argument.*

In fact, the real power of BBNs comes when we apply the rules of probability to propagate **consistently** the impact of evidence on the probabilities of uncertain outcomes. A BBN will derive **all the implications** of the beliefs that are input to it. This is how BBNs can expose some of the common fallacies in reasoning due to misunderstanding of probability.

2.4 What about the complexity of calculations on BBNs ?

Calculating all the revised probabilities once evidence is entered in a large net with many dependencies and nodes which can take on more than two values can be a very difficult task, requiring huge computational resources. From a theoretical point of view, the problem can be shown to be NP-hard.

This observation, until relatively recently, meant that BBNs could not be used to solve realistic problems. However, in the 1980s researchers discovered propagation algorithms which were effective for large classes of BBNs. With the introduction of software tools that implement these algorithms (as well as providing a graphical interface to draw the graphs and fill in the probability tables) it is now possible to use BBNs to solve complex problems without doing any calculation by hand.

For example, propagation on the 101 nodes BBN that we describe in § 5 is performed in less than one second by the tool HUGIN explorer on a PC with apentium 100 processor.

3. THE SERENE METHOD (QUICK OVERVIEW)

3.1 General principles

Like in any application domain the building process of a BBN to represent a safety argument will require three steps:

- identification of nodes: the properties (e.g. suitability of design, competence of the development team, effort spent on validation...) of the safety-related system and its development which can affect safety are identified: they will become the nodes of the BBN model,
- identification of the structure of the BBN: the relationships between these properties determine the structure of the BBN. The relationships are causal but subject to uncertainty: for example, small software modules are more likely to increase testability, but it is still possible for the testability to be high in software without modules,
- filling in of NPTs: The conditional probabilities associated with each relation must be entered to complete the model.

Then exploitation of the model can take place, by (repeatedly) entering evidence and propagating. More precisely, to use a model to argue the safety of a particular system, assessments (by means of audits, tests, measurements) must be made to determine the actual value of the properties which can be observed, in order to obtain (by propagation) revised beliefs about the unobservable properties (like safety). The greater the proportion of these properties which can be assessed, the better. Those which can not be assessed are assumed to be distributed according to the probabilities obtained by propagation. This feature introduces flexibility in the use of the model.

The probabilities can be revised each time a new piece of evidence becomes available, which means that *trends can already be estimated in the early stages of a system development.*

The aim of the SERENE method and the associated tool is to provide an efficient support in all these activities. The two following paragraphs tell how this is achieved.

3.2 Hierarchical decomposition and generic patterns

Building a BBN is not an easy task and the partners of the SERENE project had considerable difficulties in the early stages of the project, in spite of their good knowledge of the safety assessment domain.

They felt that a possibility to build a BBN in a hierarchical, top-down manner, would help very much both the construction and the understanding of the model. For example, the BBN presented in section 5 was built in two levels. This model has a rather complex structure, and 101 nodes. It is therefore impossible to represent it on a single, flat graph.

Another major finding of this project was that, in spite of the fact that BBN models of safety arguments that were produced by SERENE partners had a number of differences, reflecting the different perspectives and industry sectors of the partner companies, it was possible to extract from them a surprisingly small number (five) of typical patterns, corresponding to generic reasoning schemes.

The SERENE tool, which is currently under development, will support both the hierarchical decomposition, and an efficient mechanism to import BBN typical patterns, thus providing the user with the ability to build a BBN in a « bottom-up » approach.

3.3 Eliciting the probabilities

Another important issue in the building of a BBN, is the elicitation of NPTs. Since we are in a domain where very little formalised feedback of experience is available, we have to elicit probabilities from experts. The SERENE method provides guidance in the elicitation process, by describing the many kinds of biases experts can be subject to, and how to avoid them or at least how to reduce their impact. One simple example is the fact that one gets much better estimations in terms of frequencies of events rather than probabilities.

4. HOW EDF USUALLY DEVELOPS AND ASSESSES ITS I&C SYSTEMS

4.1 Overall organisation of the development of an I&C system

EDF usually does not develop the I&C systems on its own. In the main steps of a typical acquisition process, EDF successively:

- expresses its needs in a "system requirements specification" document,
- launches a bidding process, and selects one or, in some cases two (in order to avoid problems due to a monopoly) supplier(s),
- monitors the development conducted by the chosen supplier(s) (by peer reviews, audits, various kinds of assessments),
- performs acceptance tests on the delivered system,
- in some cases, EDF integrates parts it has developed with the supplier's parts,

- finally validates the whole system with on-site tests, before putting it in real use.

4.2 Overall organization of the safety analyses

Several teams of EDF carry out the various safety analyses and tests of the I&C system (as a whole, i.e. including the parts developed by EDF and those provided by external suppliers). Some teams focus on the evaluation of hardware, others on the analysis of system architecture and software, while further teams focus on commissioning tests covering groups of co-operating systems & equipment. One team synthesizes the results of all these tests and analyses, and delivers the global, "consolidated" safety argument to the licensing authorities.

4.3 Methods, techniques and information used to construct the safety argument

The main methods, techniques and information used to construct the safety argument are listed below:

- audits: analysis of the supplier development cycle (e.g. the organization set up by the supplier for the system development, means and techniques used, etc.), verification of the application of the supplier Quality Plan, verification of the existence of some documents produced during the system development, supplier's experience in the area.
- source code analysis: conformance with programming standards commonly used or defined in the system requirements specifications, static analysis, (e.g. control or data flow analysis), etc.
- modeling of the most sensitive parts of code, either by the role they play or by the fact they implement subtle mechanisms, like interruptions.
- review of the tests and verifications carried out by suppliers (choice of test criteria, choice of test data, etc.).
- test of the system under various environmental conditions (seismic, vibration, pressure, temperature, resistance to electromagnetic and radioactive effects, resistance to defects or variations of power supply, etc.).
- functional validation tests (test of performances, compatibility of interfaces with other I&C systems).
- documents on the feedback experience acquired by the manufacturer, in particular when some software modules are already used in other operational systems.

In order to reduce the scope of the BBN that we give in this article, in the following paragraphs, only the "functional safety" is considered. Functional safety is the ability of the system to avoid insecure behaviour in the absence of hardware failures, human errors, abnormal environmental conditions (such as vibrations, high temperature...). In other terms, functional safety is a perfect suitability of the system with respect to ***real*** *needs.*

5. A BBN FOR ASSESSING A SYSTEM DEVELOPED BY A SUPPLIER

5.1 Context

The systems considered in the BBN we are going to describe are the I&C systems used in nuclear power plants (NPPs).

In a NPP, an individual system represents only a small part of the design of the plant and of its I&C. Important issues like safety analyses or risks mitigation are not handled at the level of individual systems, but at the level of the complete power plant and its I&C. The plant level analyses and design identify the individual systems of the power plant (which may be mechanical, electrical or computer based), and provide the basis for the specification of the main requirements applicable to each individual system.

As a consequence, what EDF requires of an individual I&C system (and particularly of an individual computer based I&C system) is that its requirements specification and its specification are consistent with the inputs provided by the overall plant analyses and design, and that it complies with its own requirements and specification. Safety analyses are not considered relevant at the level of individual systems.

5.2 Methodology

We always worked as a group, through meetings, following the steps given in section 3, but without the support of the SERENE method (because it didn't exist at that time !).

In several occasions, the group could confront its ideas to those of other partners of the SERENE consortium, which caused very vivid and interesting debates. One of the first issues that were debated was the number of states that should be chosen for variables of the BBN. We decided that this number should be even, in order to force the assessor to choose between a rather positive, or a rather negative judgement, in the nodes where evidence would be input. Choosing 4 states would lead to large NPTs (64 values to be specified for a node having 3 parent nodes !). It would be impossible to justify such a number of values, and anyway, the process of evaluation of these values would be excessively long and cumbersome. Thus, it was decided to use only 2 states, meaning good or poor, acceptable or not acceptable, yes or no, etc. Moreover, in order to simplify the description of the BBN, all nodes correspond to sentences such as: *object x is y*, with the two possible values: "yes" or "no".

Example: "Requirements specification is suitable"

We also decided, after a few attempts to evaluate the values of the probabilities, to choose them (in most cases) from the following set: {0, 0.25, 0.5, 0.75, 1}. The rationale for this choice is that it would be impossible to justify a higher number of levels in this scale, which is just what is needed to express the following concepts: impossible, improbable, probable, quite probable, certain.

Another advantage is that the complement to 1 of any probability from the scale is also in the scale.

In fact, values 0 and 1 (which model deterministic knowledge) may cause some inconsistencies in the model,

which would not appear with any other value, no matter how close to 0 or 1 they would be.

For example, having declared that the probability for the state "yes" is equal to 1 for a node without parent, makes it impossible to choose the value "no" for this state in the exploitation of the BBN. This is quite normal, in the light of BBN theory, but it means that probabilities of 0 or 1 should be avoided, which is not very practical for the description of the BBN. In fact, in our BBN, values 0 and 1 as probabilities should be understood as "very close to" 0 or 1.

5.3 The BBN itself

The BBN presents an overview of how EDF assesses the processes applied by the supplier and by EDF in order to ensure / assess the functional safety of an E1B system delivered by the supplier, and how EDF integrates the view points and opinions of various technical experts. It focuses only on functional safety of systems provided to EDF by external suppliers. In particular, it does not address:

- hardware issues (random failures, resistance to aggressive environment: radiation, high temperature, vibrations...),
- human factor issues,
- global issues encompassing other systems & equipment and the overall design of the power plant and of its process,
- modification of systems already in operation,
- how technical features of the I&C system are assessed.

These limitations are justified by the necessity to limit the size and the complexity of the BBN.

Because of the general principle given above (each node corresponds to a simple sentence, with two possible states: yes or no), the description of the BBN simply consists of its topology, and, for each node, of the following elements:

- a detailed description of the meaning of the sentence,
- the node probability table,
- a rationale explaining the numbers given in the node probability table.

The BBN is *organised in two breakdown levels*. § 5.3.1 (just below) will give the description of the top level, and § 5.3.2 will describe one of the sub-nets mentioned in the top level.

5.3.1 Top level model

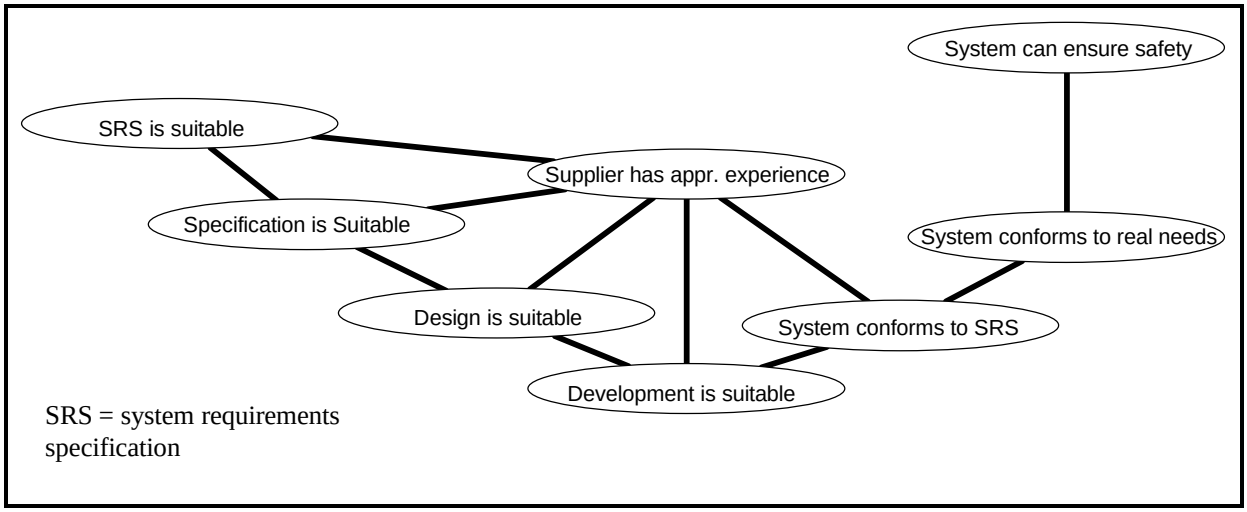

Figure 2: global structure ofthe BBN

In Figure 2, the nodes correspond to sub-nets, i.e. sub-graphs of the total BBN, and the lines correspond to the existence of communication between two sub-graphs. This communication comes through the existence of a common node, called "join node".

The global structure of the BBN is a chain that follows the stages of a development process.

Each step in this process serves as an input for the next one, and has a great influence on it.

Each step involves a *construction* process and a *verification* process, and the suitability of the final result is influenced by the appropriateness of both processes. "Appropriateness" for verifications means that the results of these verifications are *relevant*: the verifications cover enough aspects, and are enough exhaustive.

The experience of the supplier of the system is quite important, and introduces a certain degree of dependence between all stages.

The final sub-net, labelled "System can ensure safety", is reduced to a single node, and is the "output" node of the BBN.

The probabilities of the states of this node ("Y" or "N") give the level of confidence that one can have in the I&C system.

The various sub-nets are relatively similar in their structures, and in their node probability tables. The following paragraph describes a typical example: the sub-net modelling the system development phase of the life-cycle (labelled "Development is suitable" in Figure 2).

5.3.2 Detail of one of the sub-nets (Sub-net "Development is suitable")

The graph of this sub-net is given in Figure 3, at the end of the paper.

The nodes "marked" by a name beginning with "J-" (in grey) are join nodes: they belong to two different sub-nets (or more). These join nodes ensure the communication of information between the various sub-nets of the argument. The nodes marked with a dot are those for which evidence will be entered if available.

We now give two examples of node definitions, with a textual description of the ideas expressed by the corresponding NPTs:

Node "Verif. by EDF is appropriate (forDev)"

Verif. by EDF is appropriate (for Dev)	This verification assesses some key properties of the Dev: • Is complexity of developed components mastered and justified (checked by static analysis of the source code) ? • Will it lead to a maintainable system?

The corresponding NPT expresses the fact that the competence of experts is of paramount importance. In particular, the probability of an appropriate verification is set to 0 whenever the experts are not competent.

Node "Verif. by supplier is appropriate (forDev)"

Verif. by supplier is appropriate (for Dev)	This verification assesses some key properties of the Dev: • Is complexity of developed components mastered and justified (checked by static analysis of the source code) ? • Will it lead to a maintainable system? • Are the programming rules respected ?

The corresponding NPT is even more pessimistic than the previous one, because the verifications performed by the supplier necessitate a lot of tests, and therefore are very sensitive to the amount of resources.

6. EXAMPLES OF CALCULATIONS WITH THIS BBN

Here are some illustrative calculation results obtained from this BBN with the HUGIN explorer tool.

The following table gives the probabilities for the values « yes » of the main nodes of the network, i.e. the nodes corresponding to the various stages of the lifecycle. The columns of the table, numbered 1 to 5, correspond to the scenarios explained after the table.

Scenario n° ->	1	2	3	4	5
System requirements specifications are suitable	.576	.9999907	0	1	0
Specification is suitable	.419	.9999483	0	1	0
Design is suitable	.265	.9996823	0	1	0
Development is suitable	.173	.9979557	0	.9999598	0
System conforms to SRS	.138	.9952124	.113	.9983177	.243
System can ensure safety (= System conforms to real needs)	.112	.9670761	.138	.9822346	.279

1 - A priori probabilities. In the absence of any evidence, the predictions become more and more pessimistic as the development progresses. This is quite normal, since no evidence of successful verifications is entered.

2 - No evidence is entered, except the fact that all verifications give good results: this is always true for safety critical systems. Such a system can not be a candidate if it does not pass one or more tests. The probabilities are very close to 1, especially in the first stages of the life-cycle. This seems quite reasonable, since all NPTs of root nodes contain probabilities of .75 for the positive value of these nodes.

3 - Worst case calculation: negative evidence is entered in all nodes analogous to those marked with a black dot in Figure 3 (but the verifications and tests are still assumed to give good results). In this case, the probabilities are close to zero, except in the last stages, where the evidence coming from the test results gives some (little, in fact) hope.

4 - Best case calculation: positive evidence is entered in all nodes analogous to those marked with a black dot in Figure 3. The result is very much like the result of scenario 2, except that probabilities are closer to 1.

5 - Case with conflicting evidence: inputs to SRS are not suitable (the fact of supposing that this is known is very

pessimistic, which explains the zeros in column 5), all verifications give good results, but they are not exhaustive: they are performed by competent actors, using good methods, but lacking resources. This lack of resources is known, which is reflected in the BBN by setting all nodes "EDF opinion on verification activity X is good" to NO.

The results of the above simulations show that the BBN gives consistent results, which are in accordance with intuition. A "validation" of this BBN will consist in trying it on numerous scenarios, and comparing the results of calculations with the experts expectations.

It is important to note that, in spite of the fact that the results are numerical, they are nothing more than an elaborate presentation of experts **beliefs.** They must not be taken as real probabilities, and must be used only for comparative purposes.

7. CONCLUSION

Our experience is that a Bayesian Belief Network is easy to understand, and can very much improve the communication quality between experts: such a language has far more expressive power than, for example, check lists, even sophisticated ones, making use of weights. Therefore, BBNs could facilitate discussions between assessment experts.

In this paper, we have shown an example of BBN which could be used to model a safety argument, with some simulated applications.

This BBN is a prototype; its construction was a good occasion to make explicit, and capitalise the knowledge of experts in assessment of computer based safety critical systems.

This know-how will be cast, together with the know-how coming from other partners of the SERENE consortium, in the SERENE method and the associated tool.

Thanks to their help, it will become possible to build a new BBN to model each safety argument, making it much easier to understand.

A growing use of the SERENE method should help in determining which are the most influential variables in the design process of a system, which is a necessary prerequisite for setting up any kind of field experience collection.

A long term goal for the SERENE method could be to yield reliable forecasts, but to achieve this goal will require many years of calibration, by comparison between the predictions given by the method, and the real, observed safety level of the evaluated systems.

8. AKNOWLEDGEMENTS

The writing of this article would not have been possible without the fruitful collaboration between the (numerous) actors who took an active part in the SERENE consortium. We cannot quote all the individuals, but at least, here is the list of their companies, and nationalities: ERA (project manager) and the Centre for Software Reliability from UK, Objectif Technologie and EDF from France, TÜV Nord from Germany, and HUGIN from Denmark.

REFERENCES

1. P.G. BISHOP and R.E. BLOOMFIELD, "The SHIP Safety Case", *in SafeComp 95 Proceedings*, 1995

2. CASCADE, ESPRIT project n° 9032, "Generalised Assessment Method (GAM)" part 1 (rules, doc. ID: CAS/LR/WP2.T3/SM/D2.3.1) and part 2 (guidelines, doc. ID: CAS/IC/MK/D2.3.2/V3), Jan 1997.

3. CACM, "Real world applications of Bayesian Networks", *Communication of the ACM, special issue*, 38 (3), pp. 24-57, 1995.

4. PEARL J., "Probabilistic reasoning in intelligent systems: Networks of plausible inference", Morgan-Kaufmann, San Mateo, California, 1988

5. JENSEN F.V., "An introduction to Bayesian Networks", Springer Verlag, New York, 1996

BIOGRAPHIES

Marc BOUISSOU
Electricité de France,
DER/ESF Section,
1 av. du Général de Gaulle
92141 Clamart cedex.
FRANCE

E-mail : Marc.Bouissou@edfgdf.fr

Marc Bouissou has over 16 years of experience in the reliability engineering field. He has led the development of highly innovative tools, based on AI techniques, to support the activities of reliability engineering, and PSAs for nuclear power plants. His recent work is about RAMS allocation, computer controlled systems, architecture optimization. He is the vice-president of the "Methodological Research" working group (with 80 members) of the French ISDF (RAMS Institute) association.
He was awarded a degree of the "Ecole Nationale Supérieure des Mines de Paris" engineering school in 1980.

Fabrice MARTIN
Electricité de France,
ERMEL/PEL Section,
1 av. du Général de Gaulle
92141 Clamart cedex.
FRANCE

E-mail : Fabrice.Martin@edfgdf.fr

Fabrice MARTIN has over 5 years of experience in the reliability engineering field. He currently works in dependability assessment of an instrumentation and control system for energy transport and distribution. He also works on the protection plan of EDF electrical network. He specialised in computer science and real time in the « Ecole Supérieure Télécom de Bretagne » in 1991 and got a degree of the «Ecole des Arts et Métiers» engineering school in 1990.

Alain OURGHANLIAN
Electricité de France,
EP/CCC Section,
6, quai Watier
78401 CHATOU
FRANCE

E-mail : Alain.Ourghanlian@edfgdf.fr

Alain Ourghanlian has been working for EDF for 3 years. His previous experience was about the design and verification of instrumentation & control systems (I&C systems) in avionics. His present work is about new methods to specify, design and verify I&C systems in nuclear power plants. He graduated from the «Ecole Supérieure d'Electricité» engineering school in 1990.

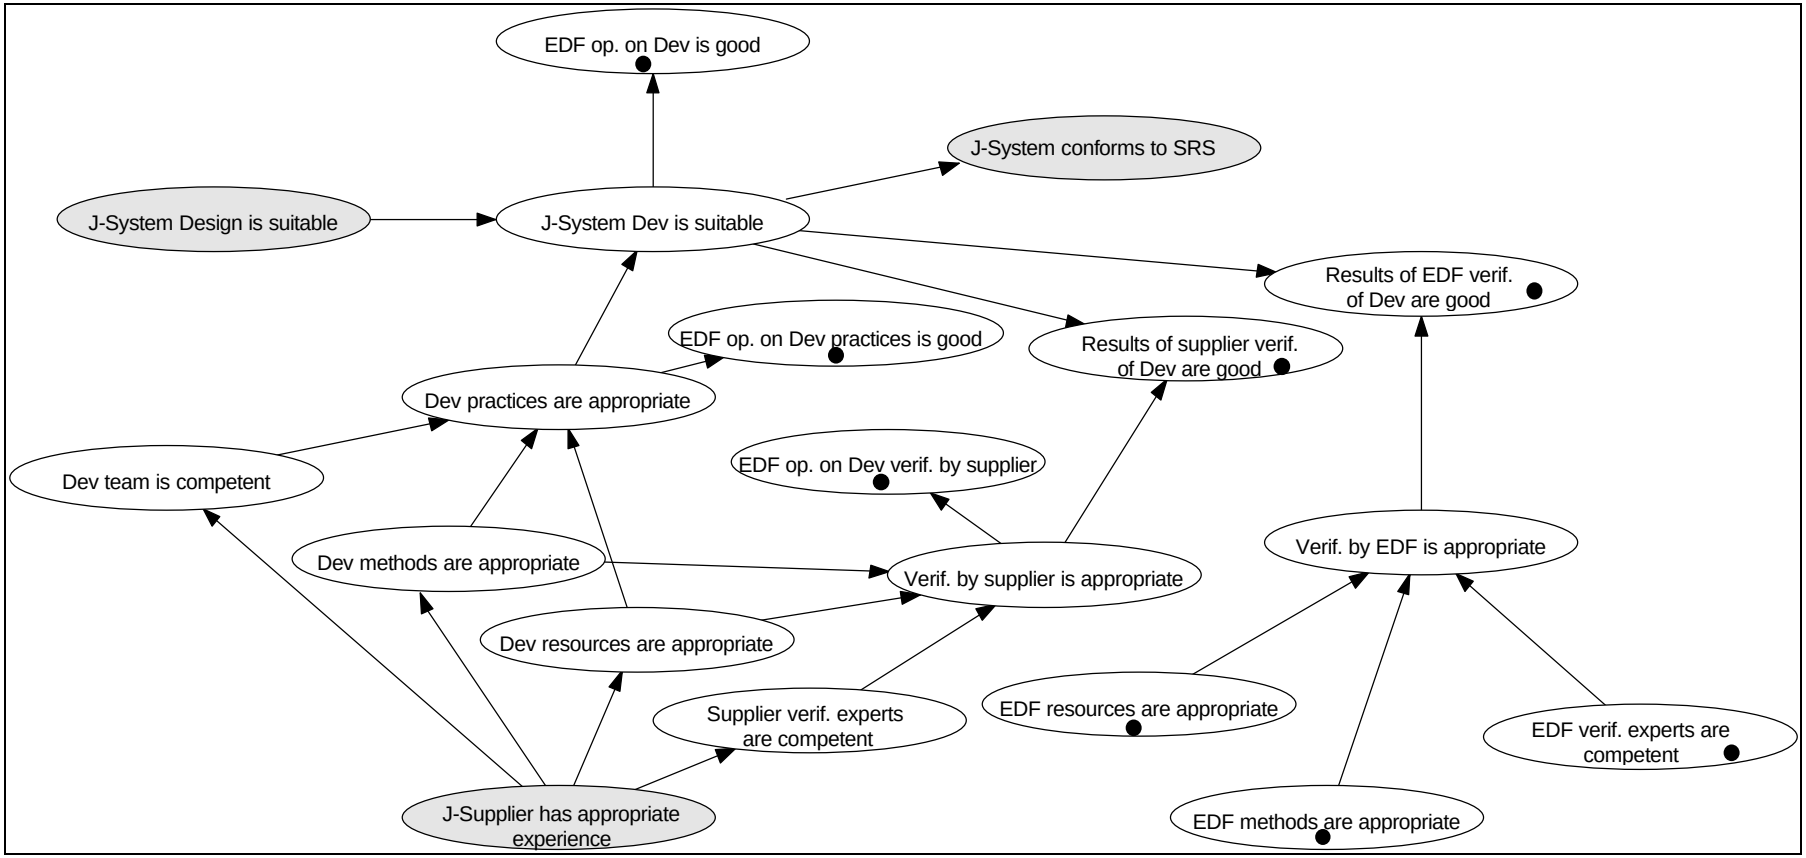

Dev : development
EDF op. on x: EDF opinion on x

Figure 3: Sub-net "Development is suitable"

(Annex of the paper "Assessment of a Safety-Critical System Including Software: A Bayesian Belief Network for Evidence Sources")