Real-time risk assessment and decision support using Bayesian networks

Andy Bolsover, DNV GL, Aberdeen, UK.

Andreas Falck, DNV GL, Oslo, Norway.

Managers and operators of major hazard facilities make complex decisions as a part of their daily work activity. These decisions are made against the background potential for a major accident.

Such decisions may be required to account for changes in a large number of factors including plant condition and performance, operational status, knowledge and experience of personnel, interactions with other activities, and the effectiveness of processes and safety systems. The information involved in the decision comes from multiple sources, may be known with varying degrees of uncertainty, and may be difficult to assess.

Conventional risk assessment methods provide a useful picture of major accident risk but some widely accepted approaches suffer from some significant problems which limit their value as tools for day-to-day operational decision making. QRAs often take the "top event" as a starting point. The ability of a QRA model to quickly assess site specific issues and identify the root causes of risk is therefore limited. Other problems may include:

- The time requirement for a technical risk assessment
- Use of generic industry-average failure data that may not reflect actual equipment condition
- The omission of competence, experience and other human factors from the assessment
- Limited account of incident/ accident history
- Difficulty in accounting for findings from inspection, test and maintenance
- Difficulty in accounting for day-to-day changes in operational status
- How to give appropriate credit for audits, safety critical system performance etc.

The paper describes DNV GL's investigations into an alternative approach that addresses many of these difficulties; it illustrates how risks can be monitored in real-time and enable safer decision making. The method is applicable to the assessment of a wide range of major accident hazard scenarios. Notably, the method can be used as the core of a complete risk assessment method which also identifies the probable causes of potential failure scenarios. The approach allows the risk profile of a major hazard facility to be updated in real time in response to new information as it is received. The speed of the approach also points to its potential in real-time detection and control systems

The method employs a Bayesian network to perform the risk assessment. Bayesian nets have been used to aid decision making in many different situations and industries, but have received relatively little attention as risk assessment and decision tools in major hazard industries. The paper will include a description of the benefits offered by this technology as well as a view of its limitations.

DNV GL acknowledges the work of Decision Systems Laboratory at the University of Pittsburgh, in developing the SMILE / GeNIe software [Ref 1] that is used to construct and solve the Bayesian network described in this paper.

Introduction

The work described in this paper was triggered by the need of clients working in a high hazard industry to make complex operational risk assessments (ORA) as part of their daily work. The paper focusses on one type of hazard in one industry – the potential for a blowout event on an offshore drilling rig – but the findings are applicable to other types of major hazard in other industries.

Many operational risk assessments are performed in response to unplanned events in which the risks at a high hazard facility are perceived to have increased. In the offshore oil and gas industry, daily variations in risk levels may be associated with events such as: the failure of an item of safety critical equipment to meet defined performance standards; lack of confidence in the competence of personnel and the potential for human error; and the ineffective operation of management procedures that are intended to maintain a safe operating environment. Other risk assessments are done as part of the planning process in order to identify the effect of proposed activities on current risk levels; offshore installation managers need to understand the cumulative effects of additional activities and identify any necessary risk management measures that may be required to allow the simultaneous operations to proceed safely.

The offshore installation manager of a drilling rig when faced with problems such as these needs to know answers to some specific questions such as: is it safe to continue operating, are remedial actions necessary, and what types of actions will be practical and effective? Furthermore the answers to these questions should be obtained quickly so that the necessary actions can be timely.

Conventional forms of ORA address these questions through a qualitative approach that may involve personal knowledge of plant status and information from inspection, test and maintenance records. Application of judgement then allows an assessment to be quickly obtained. However, conventional ORA will establish only a qualitative measure of risk. The outcome of the ORA can include the identification of risk reducing actions that may be added to a list of temporary fixes, planned inspections, and maintenance activities with no clear prioritization.

A key requirement of any operational risk assessment should include the ability to account for the cumulative effect of many risk factors that are minor by themselves but are, collectively, a cause for concern. This suggests additional requirements for the risk assessment: the assessment should be at least semi-quantified (or fully quantified) in order that the risks can be properly assessed and decisions made against defined criteria for the cumulative risk. Quantitative assessments using today's widely used methods such as event trees and fault trees can be performed but these may be too time-consuming to meet the operational needs.

Bow-ties diagrams (e.g. reference 7) are useful tools for representing major accident events and understanding the factors that affect such events. In terms of a bow-tie model, it is the author's experience that QRA studies typically focus attention on the right-hand side of the bow-tie" (which describes the consequences of a top event) while most QRA studies are relatively weak in their modelling of the left-hand side of the bow-tie (describing the causes of the top event). A reason for this weakness is that many QRA studies use generic industry average statistics as a basis for estimating failure frequencies which are then used as input data to the study. In these cases, the QRA uses top event frequencies as items of input data with no modelling of the causes of the top events. The omission of modelling of the left-hand side of the bow-tie means that the ability to quickly assess site specific issues in a QRA study is limited.

Application of Bayesian networks

Remarkably there is a method that can address these types of requirement. The method has not been widely adopted for the assessment of major accident hazards in the oil and gas industry, although it has been applied elsewhere to many different types of risk assessment and decision-making problem. The method involves Bayesian networks (BN).

References 5 and 6 are examples of the limited application of Bayesian networks within the oil and gas industry: reference 5 describes the use of BN to assess a gas explosion hazard, and reference 6 included an examination of the use of BN in well control. DNV GL is also currently employing BN to corrosion risk assessment in pipelines. Other applications in the oil and gas industry (unrelated to the assessment of major accident hazards), include optimisation of inspection planning and the evaluation of drilling prospects. A survey of the wider use of Bayesian methods identified a diverse range of applications in other fields such as: medical expert systems and diagnosis decision support systems, investment and strategic planning, social policy decisions, engineering projects, and the assessment of emerging risks associated with new technology.

The Bayesian method

Bayesian networks get their name from Thomas Bayes (1701-1761) who originated a branch of statistical and probability theory that defined the principles of probabilistic inference and a conditional probability theorem known as Bayes' rule. The necessary calculations to use Bayesian statistics tend to be complex, so Bayes' work had limited application when it was originally published (Ref. 2).

There are a large number of more-modern references which describe the Bayesian method. Many of these assume that the reader already has a strong foundation in Bayesian method, but references 3 and 4 provide easy introductions. Reference 3 contains several easily accessible papers on the application of Bayes work to decision analysis. Reference 4 describes the computer application of Bayesian method using one of the commercially available tools.

The complex calculations that were necessary to perform Bayesian calculations were significantly assisted by the development of modern computers, and academic work to develop efficient computer algorithms (e.g. reference 1) for solving Bayesian problems.

Bayesian methods explain how evidence from observations can be used to make systematic inferences about events that have uncertain probabilities of occurrence. Application of Bayes rule typically allows evidence to be used to refine initially uncertain values of probability in a way that the uncertainty is reduced.

The Bayesian method can be illustrated by a small example.

In the context of an offshore drilling operation the probability of a blowout is affected by many factors which include: the competence of people in key roles, the reliability of safety critical items of equipment, and the type of drilling operation. In the absence of any information about the drilling rig, its crew, or the characteristics of the well that is currently being drilled, we can still be confident that the probability of a blowout must be in the range 0 to 1. If we then learn that the crew is highly competent, application of Bayes' rule tells us how to reassess the probability of a blowout. If we also learn that the equipment is highly reliable, we can make a further reassessment. As more is learned the probability estimate becomes more refined until eventually we are able to say that probability of a blowout on the current well is most likely to be in the range say $1x10^{-5}$ to $5x10^{-3}$. Note that uncertainty never disappears in the Bayesian method; the blowout probability may be outside this range. Further information about specific items of equipment or specific people allows the estimate to be further adjusted.

A key concept in the Bayesian method is "probabilistic inference". Evidence about the reliability of safety critical equipment is used to "infer" the changes to the probability of the blowout. The chain of inferences may not be direct, but could involve the influence of multiple factors that affect the probability of a blowout. These influencing factors are made visible in a 'Bayesian network', a tool that is widely used in modern Bayesian probability calculations. Figure 1 shows a very simple Bayesian network which describes some of the factors (represented by the diagram *nodes*) that can influence the probability of a blowout.

The example refers to a mud engineer who must maintain a column of mud in the well bore to hold back the reservoir pressure. Various operations during drilling can result in variations of mud column density and allow the reservoir pressure to cause a "kick" which is seen at the surface. The example shows three influencing arcs and the causal directions between those influences:

- The competence of the mud engineer influences the probability of a critical error in the mud mix
- The competence of the mud engineer affects his response to a kick
- The mud mix will also affect the probability of a kick

The three tables in Figure 1 contain probability data that show how the nodes influence each other. The links (*arcs*) in the diagram show where these influences exist. The example network contains only three nodes; real practical Bayesian networks may contain several hundred nodes and influence arcs. The nodes in a Bayesian network can be in different *states* of occurrence. In this simplified example, the *states* of the "competence node" can be 'good' or 'poor', the mud mix can be 'error' or 'ok', and the response to the kick can be 'failure' or 'ok'.

Figure 1: Simple Bayesian network showing nodes, arcs, states and probability tables

Competence	
good	poor
0.8	0.2

	Critical error	
Competence	error	ok
good	0.003	0.997
poor	0.01	0.99

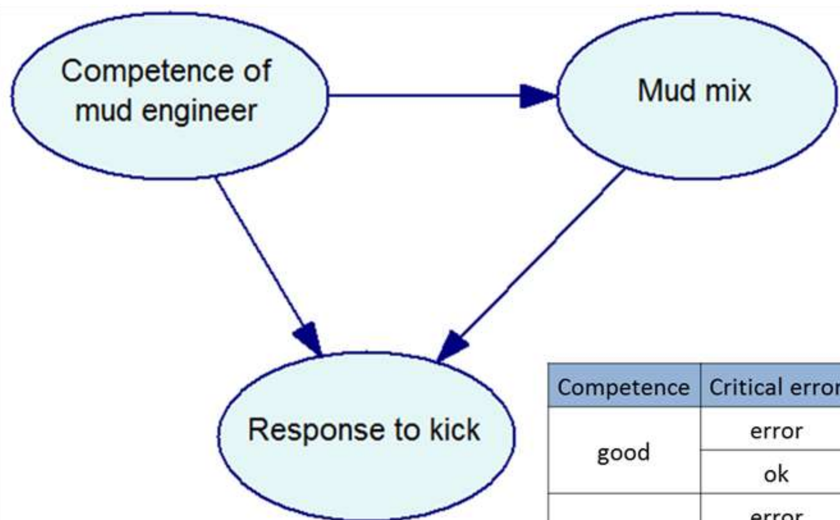

		Failure to respond	
Competence	Critical error	failure	ok
good	error	0.002	0.998
	ok	0.001	0.999
poor	error	0.01	0.99
	ok	0.005	0.995

The data describes how the influencing nodes affect the probability of occurrence of each state. In this example, the level of competence of the mud engineer is uncertain as are the probabilities of an error in the mud mix and a failure to respond to a kick in an appropriate manner. According to the data in this Bayesian net, the competence of the mud engineer is either 'good' (probability = 0.8), or 'poor' (probability = 0.2), and the probability of a critical mud mix error is conditionally dependent on the engineer's competence as shown by the table of probabilities at top right in Figure 1. Similarly the probability of failure to respond correctly to a kick is conditionally dependent on both the mud engineer's competence and the mud-mix.

In this example, each node has only two *states* e.g. the engineer's competence is either 'good' or 'poor' while in practice, there may be many intermediate grades of competence. Similarly the mud-mix and kick response could each be described by a more refined set of states. The nodes in practical full scale BNs may have many more states – sufficient to adequately describe the range of possible conditions at each node.

Powerful computer algorithms are necessary to calculate the probabilities at each node in a practical full-scale BN, but for the simple example in Figure 1, the probabilities can be calculated by hand. In this example, the probability of a mud mix error is 0.0044, and the probability of a failure to appropriately respond to a kick is about 0.0018. However, if it is known that the competence of the mud engineer is "good", the probability of a mud mix error is 0.003, and the probability of a failure to appropriately respond to a kick is about 0.001. Details of the calculations are not presented here, but they are relatively straightforward and left as an exercise for the reader.

These example calculations use probabilistic inference in the direction from causes to effects, i.e. in a forward direction, but the solution algorithms also automatically make inverse inferences, from effect to cause.

For example if it is observed that the mud engineer has failed to respond to a well kick, then it is possible to infer the revised probability that the mud engineer has "good" competence. In this case, the probability that his competence is "good" would be calculated as 0.443. This inverse calculation is rather more difficult to do by hand, but computer-based methods are highly optimised and can typically make such calculations very rapidly – even in very large realistic BNs.

This simple example could be expanded to include other relevant nodes representing factors such as the reliability of hardware, other aspects of individual reliability and the influence of management failings etc. Some of these factors are included in a larger demonstration model which is described in the next section of this paper.

SOUL risk assessment tool

A Bayesian network has been used to model the occurrence of an offshore drilling blowout event. Blowouts are prevented by the correct operation of primary and secondary well control systems which rely on a combination of technical hardware systems, drilling procedures and personnel competence. A major feature of the primary well control system is the mud column in the well bore which holds back the pressure in the reservoir, and a key element in the secondary well control system is the blowout preventer (BOP) which is required to close in the well if the primary well control system fails.

Figure 2: Bayesian network for an offshore drilling blowout

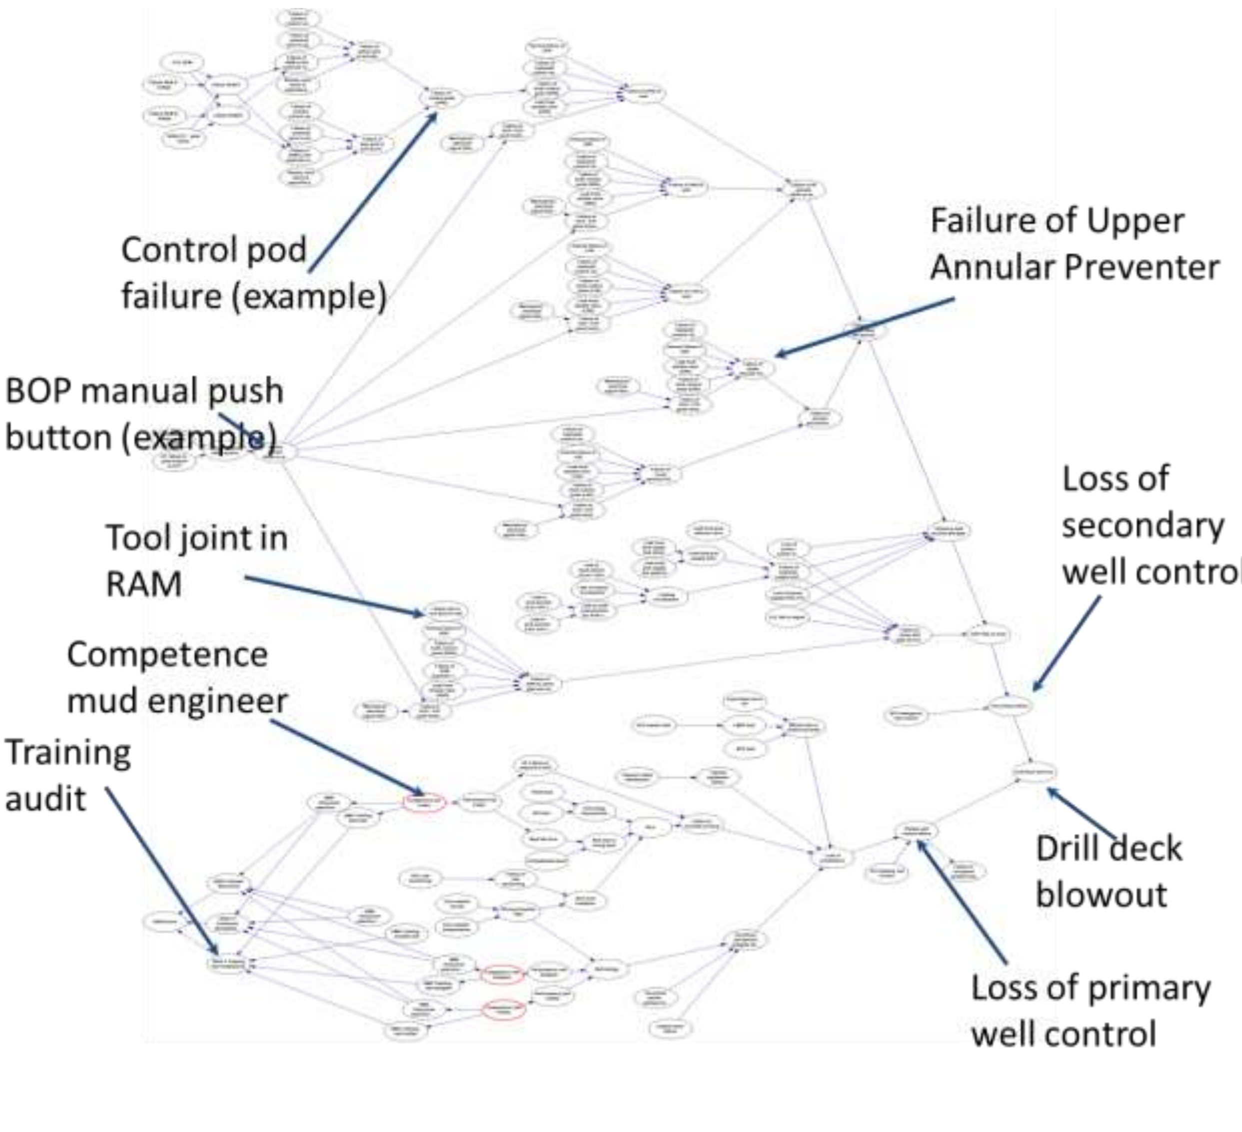

The Bayesian network has been developed as a demonstration tool in order to explore the potential of the Bayesian method for assessing the probability of a blowout and identifying the probable causes of a blowout. The BN (shown in Figure 2) accounts for many of the major features of a well control system that affect the probability of a blowout. The model includes nodes that represent the likelihood of human error in some specific tasks, and nodes that represent the findings from audits of some critical aspects of the management system. The findings of audits do not of course directly affect the probability of a blowout, rather the audit findings are influenced by factors such as the competence of the drilling personnel. The direction of the influencing arcs is therefore from the nodes representing competence to the nodes representing the findings from audits. But, as already noted, the Bayesian solution algorithms correctly handle inferences in both directions: from cause to effect, and from effect to cause. They are therefore able to make inferences about competence from the findings of audits.

A demonstration risk assessment tool, SOUL, has been built based on the network shown in Figure 2. The name SOUL stands for Safety Offshore Upstream Landscaping. The term ‘landscaping’ in this acronym reflects the radically new approach to risk assessment introduced by the Bayesian approach. A screenshot of the tool (Figure 3) displays a simplified version of the full network, but the user can, if necessary, navigate to any of the nodes in the full network. The user can enter new information at any of the nodes in order to reassess the probability of a blowout and monitor the most probable causes of a blowout event. This allows the user to identify the need for risk management action (i.e. risk reduction) and also directs the user to those factors where risk reduction effort should be focused.

Figure 3: Screenshot of SOUL risk assessment tool

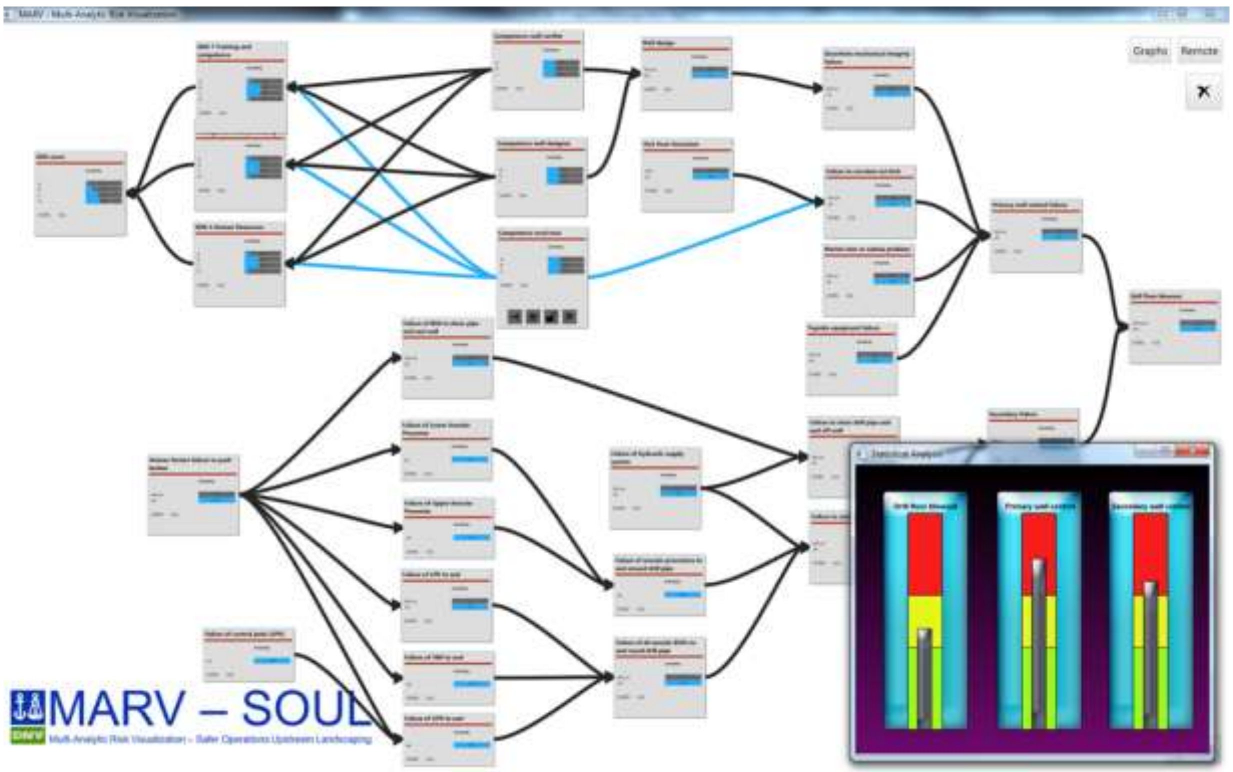

A scenario for application of the SOUL assessment tool might involve an OIM who learns that a critically important member of the crew has been replaced. The previous mud engineer was highly experienced, but the new mud engineer has relatively little experience and is not known to the OIM. The OIM’s concerns are further increased when he learns of poor performance in a recent audit of the rig operator’s recruitment, and training systems. The OIM is later told that the Upper Annular Preventer (UAP) in the BOP has failed to deliver its function in a routine test.

The OIM wants to understand the risk implications of the crew change, the poor audit finding, and the news that the UAP has failed a function test.

The OIM (or his onshore safety management support team) can assess the risk implications by using the assessment tool which includes nodes that model: the competence of personnel in critical roles, effectiveness of training, and reliability of all rams and preventers in the BOP.

The effect of the new information is shown in Figure 4. The probability of a blowout while drilling the next well increases by about 2% when the OIM sets the competence of the mud engineer to ‘low’. The OIM then sets nodes for audits of HR, contractor management, and training to “poor performance”. This has a larger effect on the probability of a blowout.

The model shows that the effect of poor performance by the mud engineer is relatively small, but the implications of poor audit performance are relatively large. How can this be? The explanation comes from the Bayesian network itself: the mud engineer affects only a few aspects of the overall well control system and the BN currently describes an effective BOP which is available to mitigate any serious failing by the mud engineer. But the implications of poor performance in the recruitment and training systems are more-widespread; the audit has raised questions about the competence of many members of the crew in addition to the quality of the well design.

The OIM then sets the UAP node to the 'failed' state. This results in another large increase in the assessment of blowout probability.

Figure 4: OIM's reassessment of blowout probability

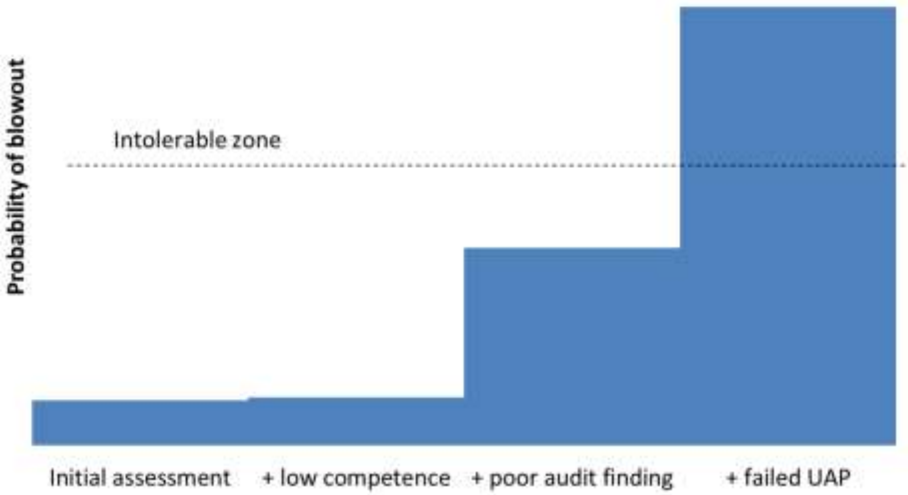

Figure 4 shows that the probability of a blowout is now intolerable. The assessment should therefore trigger essential actions.

Further investigation of the Bayesian network can reveal the most probable causes of UAP failure and can start the process of investigating the failure of this critical equipment item. This BN can also identify the relevant probabilities for all of the modelled causes of failure. The relative importance of competence and the UAP function failure can also be seen in the context of other potential contributing causes of a blowout.

Benefits and limitations of Bayesian networks for major hazard risk assessment

It is important that the strengths and weakness of Bayesian networks are properly understood; they are not a cure-all for difficult risk assessments. Important issues of limitation that may arise during the construction of a Bayesian network tool include:

- **Development time**. BNs are generally not available "off the shelf"; development of applications requires the capture of know-how and experience from subject matter experts.
- **Probability calculators**. BNs perform calculations of only probability; they do not perform calculations of consequence, frequency, risk, or even simple non-probabilistic arithmetic such as addition and subtraction. If these types of calculation are required they must be done outside the BN.
- **Thought processes**. BNs use techniques from a branch of probability and statistics that is not widely used. The "Bayesian" approach can be difficult to understand for people who are accustomed to the more-familiar "frequentist" approach.
- **BN construction**. If fast run times are to be achieved then the network designer needs to work within limitations on the design of each node. The most important limitations include: use of discrete states; small number of states in each node; and small number of incoming arcs at each node.
- **Time dependency**. Modelling of time dependency requires a step-wise treatment of time, and the number of time steps should be minimised if fast run times are required. Loops of nodes are prohibited.

None of these limitations is a significant barrier to the use of BNs. The first "limitation" may not even be a real limitation because the explicit inclusion of expert knowledge will be judged a significant advantage by many. Each of the other limitations can be handled by appropriate training and experience in the use of BNs. Careful design can mitigate the difficulties associated with construction and the treatment of time dependency.

The second "limitation" arises from the fundamental nature of Bayesian methods; Bayesian networks are powerful tools for the assessment of situations involving complex conditional probabilities. The calculation power of Bayesian networks is however entirely limited to probabilities. Bayesian networks do not calculate the frequency of failure. This means that the failure rates from the standard tables must be multiplied by the framework period of operation in order to obtain the probability of failure. As an example, the SOUL model calculates the probability of a blowout while drilling a single well.

Likewise, Bayesian networks do not model consequences. Consequences are represented in a BN by the network's nodes, and the different severities of consequence are represented by different states within the node. Where severity calculations are required, these must be done separately and used to calibrate the meaning the nodes states Very few such calibration calculations were done for purpose of constructing the SOUL model. For example, the end node in the model titled "Drill floor blowout" has only two states "blowout" and "not" i.e. a blowout either occurs or it does not; there is no refinement on the severity of the blowout. Other simple Bayesian networks modelling blowouts might include states such as and "kick" and "well release", but these are omitted in the SOUL demonstrator.

Basic arithmetic operations cannot be performed within a Bayesian network. As an example, this means that a fatality probability cannot be readily converted to a PLL within the BN although there are some "trick" methods that can sometimes do this work in a simplistic manner. In practice, a full risk assessment tool is likely to have the BN as a part of the risk assessment tool, with additional functionality outside the BN to perform basic data preparation and presentation of results as risks.

Bayesian networks have several features which provide unique advantages over some alternative risk assessment tools. The following paragraphs highlight some of the special features of BNs that make them attractive to risk analysts and describe how some limitations can be handled and allow BNs to be used as the core of rapid risk assessment decision support tool. The paragraphs also compare BNs against two alternative approaches: a technical QRA and experience-based ORA.

Speed

If a Bayesian network is well designed then it can perform reassessments very quickly. The speed of the assessment depends on the size and complexity of the network. There are various algorithms for solving Bayesian networks; the choice of algorithm results in trade-off between memory requirements, speed and precision. SOUL uses the fastest known exact algorithm, the 'clustering' algorithm (Ref. 2) and in this case the model update is completed within a fraction of a second meaning that the reassessment is essentially completed in real time. This suggests the potential application of Bayesian networks in real time control systems.

Indeed, BNs have been used in spacecraft control systems to provide real time control responses. A recent project run from NTNU (reference 6) has also investigated the potential of BNs to provide real time control of drilling system. A conventional ORA can (in some cases) be performed rapidly, but the timescale to turnaround a QRA is typically much slower and the end result is unlikely to to account for the range of factors that could be included in a BN.

Incomplete data

Bayesian networks do not require a full set of input data in order to work; they can generate results if information is incomplete. This characteristic is an extension of their ability to handle uncertainty, but the practical implications are slightly different. For example, if nothing is known about the competence of the mud engineer, this this can be omitted from the assessment and it will be performed while taking account of the uncertainty associated with this missing information. QRA systems usually require a full set of data to run although default data can be provided for some parts of a QRA model.

Prediction and diagnosis

When the input data to a Bayesian network is altered, the probabilities in nodes across entire network are updated. This means that if the reliability of the Blowout Preventer (BOP) is updated by newly received information the probability of a blowout occurring will be automatically reassessed. The nodes corresponding to the causes of change in the BOP reliability will also be updated. This means that Bayesian networks will predict the probabilities of the effects of new information, but also diagnose the probable causes of the new information. Bayesian networks therefore provide a special capability to combine prediction and diagnosis.

Handling of uncertainty

A core feature of the Bayesian method is handling of uncertainty. In this context, uncertainty does not simply mean that BNs perform probability calculations, but they handle the uncertainties within those probabilities. This means that BNs are suitable for the assessment of problems where uncertainty is inherent.

The effectiveness of organisational procedures is inherently uncertain because of difficulty of measuring the effectiveness of those procedures in any circumstance (as well as the wide and uncertain variety of complex circumstances in which they might be applied). Human error rates and competence are also highly uncertain; here the uncertainty arises from issues such as difficulty of measurement and the natural variability of human behaviour in different circumstances. Some sources of uncertainty can be directly included in a Bayesian model, but other uncertainties can be represented by the breadth of the distribution of probabilities at the node.

The risk assessment will then take account of the OIM's judgements on procedural effectiveness and personnel competence. QRA makes explicit calculations using probabilities, but the uncertainties in those probabilities are not usually represented, except through sensitivity studies. In practice some uncertainties will only be recognized through conservatisms in the assessment.

Complexity

It has already been noted that BNs are tools for handling complex conditional probabilities. Figure 2 shows an example of how the risk influencing factors can have complex dependencies, although the relationships between the nodes can be significantly more complex than is shown in Figure 2 with networks having many more nodes and cross-linked arcs. This capability can be both an advantage and a disadvantage; the tool allows the modelling of situation with multiple interrelationships between the influencing factors, but models can become unmanageably complicated. Good design of large BNs requires a careful disciplined approach. Checking a BN can be a relatively straightforward, albeit lengthy, process because each node is directly influenced only by its input nodes. Checks should therefore separately consider each node and its probabilistic dependency on input nodes. Most conventional QRAs make some use of either event trees or fault trees; these provide some of the capabilities of Bayesian networks, but are generally not able to handle the heavy cross-linking that can be found in Bayesian networks.

Conclusions

A number of conclusions can be drawn from this work:

The benefits of Bayesian networks as a tool to improve QRA are significant, but these benefits have not been widely recognised for the management of major accident hazards in the oil and gas industry. Characteristics of Bayesian networks that can assist current approaches to QRA include:

- Modelling of a wide scope of risk influencing factors such as hardware reliability, procedural effectiveness and human error rates
- Handling of uncertainty
- Speed of reassessment
- Ability to perform both prediction and diagnosis
- Handling of complex dependencies between risk factors

Bayesian networks do not directly address all the elements of a QRA. In particular, they handle only probability calculations. The interpretation of probability results from a Bayesian network to obtain results in terms of frequency is however usually straightforward and requires simple scaling. Consequence calculations must be performed separately and may be required to properly calibrate the meaning of some nodes.

In terms of the familiar bow-tie representation of major accident incidents, it is recognised that most of today’s QRA studies focus attention on the “right side of the bow-tie” and pay relatively little attention to the left-hand side. Many QRA studies are therefore not well-suited to the identification of root causes. The ability of QRAs to quickly assess site specific issues is therefore limited. Bayesian networks provide a way of addressing both sides of the bow-tie diagram. This characteristic opens the possibility of using BNs to extend QRA capabilities into the identification of the causes of failure, and the speed of BN solution engines makes then well-suited as a core element of rapid assessment tools for operational risk assessment.

References

[1] GeNIe and SMILE http://genie.sis.pitt.edu/ (accessed 9 December 2013)

[2] Bayes, T., "An Essay towards solving a Problem in the Doctrine of Chances. By the late Rev. Mr. Bayes, communicated by Mr. Price, in a letter to John Canton, M. A. and F. R. S.," Phil. Trans. Royal Society London 53: 269–71, 1763

[3] Readings on the principles and applications of decision analysis (Volumes I and II), edited by Ronald A. Howard and James E. Matheson, Strategic Decisions Group, 1989

[4] Risk Assessment and Decision Analysis with Bayesian Networks, Norman Fenton and Martin Neil, Queen Mary University of London and Agena Ltd, 2012

[5] Decision-making to treat an explosion hazard, Andy Bolsover, DNV UK and Dr Martin Wheeler, Elf Exploration UK plc, Safety on Offshore Installations, Conference Proceedings, 1999, ERA Report 99-0808, ISBN 0 7008 0700 4

[6] CODIO (Collaborative Decision Making in Integrated Operations), http://www.computas.com/en/Services/Integrated-Operations1/Codio (accessed 9 December 2013)

[7] BowTieXP, http://www.governors.nl/bowtiexp.html (accessed December 2013)