

Using Bayesian networks to model the operational risk to information technology infrastructure in financial institutions

Martin Neil¹

Professor of Computer Science and Statistics,
Queen Mary, University of London, and Chief
Technology Officer, Agena Ltd

David Marquez

Researcher, Department of Computer Science,
Queen Mary, University of London

Norman Fenton

Professor of Computer Science, Queen Mary,
University of London, and CEO, Agena Ltd

Abstract

This paper describes the use of Bayesian networks (BNs) to model the operational risk to information technology (IT) infrastructure in financial and other institutions. We describe a methodology for modeling financial losses that might result from operational risk scenarios involving data centers and operational locations, applications and systems, processes, and ultimately IT supported customer-facing services. We focus on modeling the causes and effects of unexpected loss events using a Bayesian network model of the IT infrastructure combined with assessments of the severity of impact of these events in terms of the Value at Risk (VaR) for the organization. We use a state-of-the-art Bayesian network tool to simulate an example analysis of the model. The work also illustrates how ideas commonly used to measure risk in other industries, especially the Aviation and Nuclear sectors, readily translate to operational risk in finance.

¹ The authors would like to thank David Hager, University of Stavanger, for his comments and feedback.

Using Bayesian networks to model the operational risk to information technology infrastructure in financial institutions

The Basel Committee on Banking Supervision, in reaction to a number of well-publicized financial disasters, has mandated a system of regulation addressing the issue of operational risk (OpRisk) and its assessment [Basel (2006)]. Key to the regulatory process is the need for businesses to model their operational risks, in terms of a variety of loss event types, including IT and systems failures, in order to arrive at an appropriate regulatory capital charge.

Of course, the OpRisk problem is not peculiar to the financial sector and operational risk is not a new topic. In his book, James Reason argues that operational risk is faced by all organizations and he uses examples from the financial, rail transport, civil aviation, and nuclear power sectors to support his case [Reason (1997)]. Reason identifies a host of reasons why catastrophic failures occur in these safety critical industries, including (but not restricted to): a failure to enforce lessons learnt from previous failures, slow degradation or collapse of safety procedures, changes in culture and management, lack of visibility and support for risk reporting, and lack of attention to detail. The key conclusion from this is that accidents are not solely the result of human fallibility but are supported by organizational features that fail to defend against all-too-human mistakes, slips, and (in the case of fraud) malicious acts. From this we can conclude that OpRisk prediction is inextricably entwined with good management practice and that its measurement can only meaningfully be done if the effectiveness of risk and controls processes is regularly assessed. This contrasts sharply with the view that modeling OpRisk simply involves the investigation of statistical phenomena.

By the same arguments, financial catastrophes are not a 'bolt out of the blue,' nor are they inexplicable. Financial scandals, such as Société Générale [The Times (2008)], Barings [Rawnsley (1995)], and the Allied Irish Bank [Wachtell et al. (2002)] were all the result of fraudulent activities building up over lengthy periods of time during which active management could have discovered and prevented them. Indeed, if caught early the events would not have been catastrophes

at all. There is a tendency to see financial disasters as single 'ultra high loss' events rather than aggregations of smaller losses accrued over a period of time. This is understandable given the fact that the losses have to be realized upon discovery, all at once. But this does not change the fact that such losses are accumulated daily and could be detected by good diligence, applied routinely. It is precisely this routine attention to good practice that, just as in safety critical industries, prevents disasters occurring. Any OpRisk scheme should, therefore, focus on detecting near misses and small losses on a monthly or quarterly basis before they become large losses and disasters (Jérôme Kerviel's alleged fraud at Société Générale in 2008 was said to have heavily contributed to the global financial instability in January 2008).

In this paper we argue that Bayesian networks (BNs) provide an attractive solution to the problems identified above and show how we can apply them to the problem of modeling IT failures and their implications on business services. BNs have the advantage that they enable us to combine any statistical data that is available with qualitative data and in a way that mirrors the causal structure underlying the process itself, thus making it easy to understand and communicate to business users. Using BNs we can: combine proactive loss indicators, related to the business process, with reactive outcome measures such as near miss and loss data; incorporate expert judgments about the contribution that qualitative estimates can make to the overall risk assessment; enter incomplete evidence and still obtain meaningful predictions; perform powerful 'what-if?' analyses to test the sensitivity of conclusions; obtain a visual reasoning tool and a major documentation aid; perform back-to-back comparison of alternative scenarios and sensitivity analyses for the purposes of assessing the impact of design changes to the infrastructure; provide a VaR assessment for each service and in aggregate in order to determine insurance premiums (or indeed decide to self insure) as well as determine levels of and areas for investment in improvements; and obtain outputs in the form of verifiable predictions against actual performance measures and loss event rates.

Using Bayesian networks to model the operational risk to information technology infrastructure in financial institutions

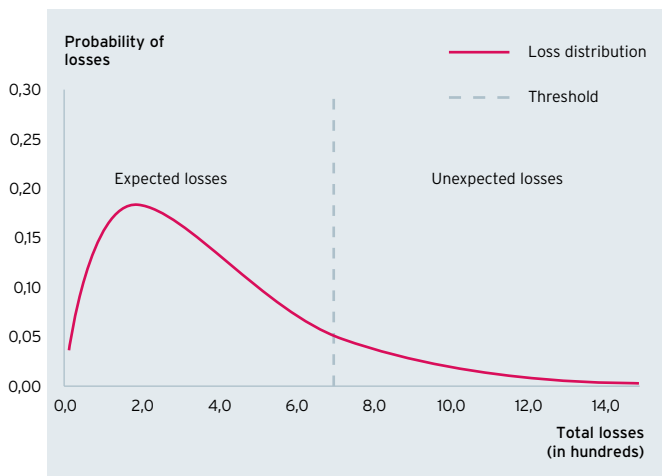


Figure 1 - Expected versus unexpected total losses

Estimating expected and unexpected losses

The Basel report [Basel (2006)] classifies financial losses due to operational factors into two 'types': expected losses, which are considered the 'normal' losses that occur frequently, as part of everyday business, with a low severity (examples include losses due to accidentally miscalculated foreign exchange transactions); and unexpected losses, which are the unusual losses that occur rarely and have a high severity (examples include losses resulting from a major fraud activity).

Figure 1 shows the distinction between expected and unexpected losses. The demarcation line is purely arbitrary (in Figure 1 this separation is shown at total losses of U.S.\$700,000). It, therefore, makes little sense to use fundamentally different methods for predicting expected and unexpected losses; it is better to think in terms of finding a distribution whose tail represents the unexpected losses.

The traditional statistical approach to these kinds of problems is to rely purely on historical data to find the inherent distribution of losses. However, in the case of operational loss data, even when a large loss dataset is available, it is

unlikely that there will be enough data on the large unexpected losses for us to be able to estimate the tail of the distribution properly – usually we end up with tails that are too thin or indeed too fat if the loss data are not relevant for the domain in question. Even when modeling the expected losses (the bulk of the distribution), the data-oriented approach suffers from the following limitations: (1) loss data will be gathered over a period of time that may represent varying levels of operational effectiveness and risk/threat (we cannot expect that losses are generated from one single distribution with a small number of known parameters); (2) losses experienced are simply a sample of possible events (they may not be representative of changing operational processes. As the underlying operational process degrades or improves the value of such historical data wanes); and (3) the reported loss data might be wrong (under-reporting and data ambiguity can lead to significant errors in estimation).




Figure 2 - Bayesian IT operational risk asset model

Using Bayesian networks to model the operational risk to information technology infrastructure in financial institutions

Any attempt to bolster loss data with data gathered from other organizations is subject to the same problems and more because very often the provenance of the data is unknown or in doubt.

Bayesian networks

A Bayesian network (BN) is a directed acyclic graph, such as the one shown in Figure 2, whose nodes represent the uncertain variables of interest and whose edges are the causal or influential links between the variables. Associated with each node is a node probability table (NPT), a statistical distribution, or parameterized function. In the case of an NPT the relationship is governed by a set of conditional probability values that model the uncertain relationship between the node and its parents together with any uncertainty that is present in that relationship.

The underlying theory of BNs combines Bayesian probability theory and the notion of conditional independence to represent dependencies between variables [Pearl (1986), Spiegelhalter and Cowell (1992)]. To date BNs have proven useful in many areas of application, such as medical expert systems, diagnosis of failures, pattern matching, speech recognition, and, more relevantly for the OpRisk community, risk assessment of complex systems in high stakes environments [Neil et al. (2001), Neil et al. (2003), Fenton et al. (2004)], including financial institutions [Neil et al. (2005)].

BNs enable reasoning under uncertainty and combine the advantages of an intuitive visual representation with a sound mathematical basis in Bayesian probability. With BNs, it is possible to articulate dependencies between different variables and to propagate consistently the impact of evidence on the probabilities of uncertain outcomes.

The key to the successful design of a BN model is the meaningful decomposition of a problem domain into a set of causal or conditional propositions about the domain. Rather than ask an expert for the full joint probability distribution of all the variables of interest, which is obviously a very difficult

task, we can apply a 'divide and conquer' approach and ask for partial specifications of the model that are themselves meaningful in the experts' domain. In our case, for IT operational risks the structure is an obvious artifact derivable from the model, as will become evident in later discussion.

Next, we need to model the NPT for each variable (node): this can either be done using historical data (including, for example, using standard Bayesian parameter learning approaches or Monte Carlo simulations) or by simply asking the expert to provide a series of subjective estimates. Ideally we would expect these estimates to be based on experience and knowledge rather than blind guesswork.

We can embed continuous and discrete statistical distributions within the BN model, as NPTs, and generate values for these NPTs by approximation methods, including Monte Carlo simulation. Until very recently, BN tools were unable to handle non-Gaussian continuous variables, and so such variables had to be discretized manually, with inevitable loss of accuracy. However, a breakthrough dynamic discretization algorithm [presented in Neil et al. (2007)] has now been implemented in a software tool [AgenaRisk (2008)], which provides an approximate solution for classical Bayesian statistical problems, involving continuous variables, as well as hybrid problems involving both discrete and continuous variables.

Figure 2 shows an example Bayesian IT operational risk model that we will discuss in more detail in the next section. Here it suffices to point out that the nodes represent processes/events/risks and the arcs represent causal/functional/physical dependencies between them.

Once a BN is built, it can be executed using an appropriate propagation algorithm, such as the Junction Tree algorithm [Jensen (1996)]. This involves calculating the joint probability table for the model from the BN's conditional probability structure in a computationally efficient manner. This is achieved by automatically deriving from the BN an intermediate graph theoretic representation of the BN,

Using Bayesian networks to model the operational risk to information technology infrastructure in financial institutions

called the Junction Tree (JT). The JT allows localized, modular computations to be executed using a message-passing algorithm. This is, in essence, an elaborate form of Bayes' theorem [Jensen (1996), Lauritzen and Spiegelhalter (1988), Pearl 1986), Spiegelhalter and Cowell (1992)]. This process is entirely automatic and, in a tool like AgenaRisk, is hidden from the domain expert. When the BN is executed the effects of data entered into one or more nodes can be propagated throughout the BN, in any direction, and the marginal distributions of all nodes updated. This makes it ideal for 'what if?' and scenario analysis.

Using BNs to model operational risk of IT infrastructures

Motivation

When modeling operational risk of IT infrastructures we must first identify what we want to model. The subjects of the model will be the interacting and interrelated IT oriented tasks and processes that together support the financial service functions of the organization. A good place to start is to use IT management processes as defined by ITIL [ITIL (2004)]. These tasks and processes are often pursued as disparate activities by different stakeholders with overlapping responsibilities and in many organizations are seen as different separate tasks. These include:

- **IT infrastructure management** – what is the reliability provided by applications, processes, and technologies that together provide the business service?
- **IT architecture planning and design** – how do we design an optimized and robust IT architecture that delivers a reliable service?
- **Business risk analysis** – can we mitigate, control, or inhibit the likelihood of risky events or ensure that their economic or reputational impact on service provision is minimized?
- **Business continuity management** – how do we recover from large-scale incidents and ensure continuity of service?

Given that IT operational risk assessment covers all of these areas, it makes sense to take a holistic view of the problem and consider the ways in which these different activities interact and how the various service quality, risk, and economic impact metrics can be usefully combined together to forecast operational risk in a meaningful way. The aim here is to provide a model that unifies each of these perspectives and can be used by each of these activities to deal with risk and uncertainty.

Asset model

To help achieve these aims we model the various asset classes in the IT architecture and the business processes that rely on that architecture to deliver business services. This involves making clear the dependencies between services, their constituent business processes, and the IT applications that help support or deliver those processes. It is helpful to think of these entities as being organized in layers, with the base layer being formed by physical buildings etc., the next would include hardware and network systems, and then an application layer, followed by processes, and then finally the services. Of course, such a layered division is not unique. We will call this the asset model and Figure 2 is the example we will use here.

Figure 2 comprises a number of asset layers: location, systems, processes, and services. There are two locations: data centre alpha and data centre beta. Data centre alpha hosts system A, system B, and primary system C. So, if the data centre fails, perhaps due to flood or terrorist attack (not shown but easily added to the model), the systems based in the data centre might also fail. If data centre beta is operational then the backup for system C will be operational and system C will be operational and able to support process Y. Unfortunately, because process Y depends on system B and system C working together, it will fail because system B has failed. And so on.

Of course we could replace each of the layers with different asset classes (such as applications rather than systems) or

Using Bayesian networks to model the operational risk to information technology infrastructure in financial institutions

indeed add other layers (such as those needed to cover physical and security hazards). The final layer in Figure 2 adds the VaR models in order to quantify the effects of failure on operational losses. When armed with this asset model we can clearly and visually determine chokepoints, single points of failure and overloads and, perhaps more importantly, make explicit the link between the economic VaR and the possible sources of risk.

Asset reliability prediction

Once we have built the asset model we can predict the reliability of the service from its constituent inputs and answer questions like “if application X and application Y fails, does the service fail?” To help answer these questions we aim to identify: common causes of reliability problems (for example, co-location of applications in building subject to flood risk); the dependency logic dictating how they interact, which can be expressed as Boolean logic using and, or, not functions; and responses, repairs, and adaptations (controls, mitigants, inhibitors, etc.) that are considered in response to events to put them right and whether those responses are timely enough to prevent loss.

We can illustrate the above by zooming in on the asset model example in Figure 2 to concentrate on process X. The marginal probability distribution for process X and its sur-

rounding nodes is shown in Figure 3. We can see that process X fails if either of systems A, B, and primary system C fails. Based on this the probability of process X failing is 11.546%. Note that this calculation takes into account the fact that systems A and B could fail because they both belong to the same data centre which could fail and take both of them down. Assuming independence between the parent systems, the probability of process X failing would be higher at 13.3%. The AgenaRisk software automatically and exactly calculates all of the possible combinations of failure and success events throughout the model from the probability values supplied in the node probability tables and the Boolean logic relationships declared by the user.

Estimating VaR

Armed with the asset model we can make decisions on the basis of clear financial criteria, such as VaR. VaR is simply a measure of risk appetite and provides answers to questions like: “what is the risk of ruin in the next year?”, “what is the expected (average) loss this month?” or “I want to be 80% confident that the losses will not exceed U.S.\$10m, what do I need to improve?” We, therefore, propose that impact assessment be carried out using VaR concepts based on a valuation of each business service’s current or future potential revenue streams. Only when armed with a keen appreciation of what is at stake can we make investment decisions about business continuity plans or design decisions about a more robust IT architecture.

Once we have estimated the probabilities of failure at the service level, we can now estimate the economic losses for each service and, in total, for the business. In our example asset model we use simple loss metrics for each of services A or B with the distribution of the losses assumed lognormal under the condition of a service failure (this seems a good choice because of the potential long tail nature of operational losses). In practice the distributions can, of course, take any functional form. However we also need to model the situation where a service does not fail, which then incurs zero losses.

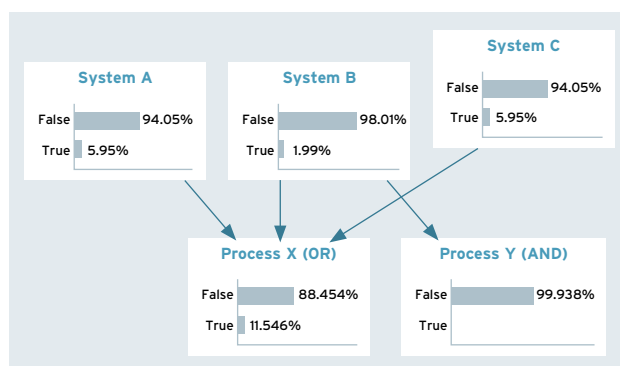


Figure 3 - Reliability quantification of the asset model

Using Bayesian networks to model the operational risk to information technology infrastructure in financial institutions

From a statistical modeling perspective severity of service failure, L , is simply calculated as a mixture under a discrete condition of failure or no failure, $S = \{\text{true}, \text{false}\}$. This can be done analytically [Venkataraman (1997)], but here we will calculate the results approximately using the dynamic discretization approach mentioned earlier. Thus, formally we have: $f(L | S = \text{true}) = \text{lognormal}(\mu, \sigma^2)$, $f(L | S = \text{false}) = 0$.

For this example we assume two services, service A and service B, whose conditional loss distributions under failure are both lognormally distributed with parameters Lognormal ($\mu_A = 3$, $\sigma_A^2 = 3$), lognormal ($\mu_B = 12$, $\sigma_B^2 = 658$), respectively.

Running a scenario

Here we run a hypothetical scenario on the model under a number of assumptions about the probability of failure of each of the different assets in the Bayesian IT operational risk model. Firstly we look at a status-quo scenario involving the currently estimated probabilities of failure and then compare this with two 'what if?' scenarios: one where data centre alpha fails and the second where data centre beta fails. In each scenario we estimate the mean loss and the VaR at 99.9%.

Figure 4 shows the summary statistics for total losses for each scenario. Notice that VaR for data centre alpha has by far the biggest impact on both service A and service B than data centre beta, with a VaR of U.S.\$92.826m as opposed to U.S.\$79.271m. If this uplift in VaR is intolerable we could then make a decision. We could decide to implement more resilient business continuity and recovery processes into the IT architecture or make the delivery services more robust to failure. Or we could decide to self insure by setting capital aside or seek insurance. The effects of either of these mitigating actions on the predicted losses can be evaluated through

Scenario	Mean loss	VaR (99.9%)
Status-quo	\$1,269,500	\$76,225,000
Data center alpha failure	\$2,359,600	\$92,826,000
Data center beta failure	\$1,405,900	\$79,271,000

Figure 4 - Summary VaR and mean losses for each scenario

comparison of their respective effect on the VaR. Comparing the predicted reduction in VaR with the investment cost of each mitigating action is useful information when deciding which course of action to take.

In addition to the summary statistics contained in Figure 4 the BN model also produces full marginal probability distribution results for all discrete and continuous variables in the model. An example of this is shown in Figure 5.

Clearly, this example model is by necessity relatively simple. However, it is still powerful and in practice, by using a tool such as AgenaRisk, the approach scales up and offers a realistic level of resolution in two ways:

- The user can specify a level of desired accuracy, constrained only by the computational resources available, and thus calculate very accurate percentile values for VaR and aggregated loss distributions over many processes.
- Use modular notation to compose large models from smaller fragments and use this to automatically generate models from databases containing loss data and process descriptions.




Figure 5 - Loss distributions for each of the services, including the total loss

Using Bayesian networks to model the operational risk to information technology infrastructure in financial institutions

Conclusion

Bayesian Networks can successfully model dependencies between events and processes in complex systems, including IT infrastructures. We have shown that they do so in a graphical way that naturally supports robust statistical and probabilistic specification of the risks and subsequent calculation of consequential risk measures, such as VaR. As such they meet the requirements of the Basel Accord [Basel (2006)], for an advanced measurement approach (AMA). Adopting a BN-based approach should, therefore, lead to better operational risk governance and a reduced regulatory capital charge. The BN approach presented here strongly contrasts with a purely statistical approach based on historical loss data alone. We believe that traditional statistical analysis techniques will neither provide good predictions of future operational risk losses, nor provide a mechanism for controlling and monitoring such losses. This second goal is obviously of utmost importance in practice.

This paper has provided an introductory flavor of BN technology and its contribution to operational risk modeling. The latest BN tools enable more sophisticated modeling and analysis than can be covered here. For example, such tools now support learning from loss data in a way that competes with techniques that have hitherto been recognized as state of the art, such as MCMC [Gelman et al. (2004)]. The new tools also make it relatively easy to model more complex dependencies between assets and losses than those covered here including, time-based losses, time to respond, time to failure modeling, and complex mixture and classification modeling to exploit consortia sourced loss data.

References

- AgenaRisk, 2008, Agenarisk Bayesian network and simulation software, www.agenarisk.com
- Basel Committee on Banking Supervision, 2006, "International convergence of capital measurement and capital standards"
- Fenton, N. E., W. Marsh, M. Neil, P. Cates, S. Forey, and M. Tailor, 2004, "Making resource decisions for software projects," in proceedings of 26th International Conference on Software Engineering (ICSE 2004), Edinburgh, United Kingdom, May, IEEE Computer Society 2004, ISBN 0-7695-2163-0, 397-406
- Gelman A., J. B. Carlin, H. S. Stern, and D. B. Rubin, 2004, Bayesian data analysis (2nd Edition), Chapman and Hall
- Information Technology Infrastructure Library (ITIL), 2004, IT service management: an introduction based on ITIL (Information Technology Infrastructure Library). The IT service management forum, Van Haren Publishing
- Jensen, F. V., 1996, An introduction to Bayesian networks, UCL Press, London
- Lauritzen, S. L., and D. J. Spiegelhalter, 1998, "Local computations with probabilities on graphical structures and their application to expert systems (with discussion)," *Journal of the Royal Statistical Society Series B*, 50:2, 157-224
- Neil, M., N. E. Fenton, S. Forey, and R. Harris, 2001, "Using Bayesian belief networks to predict the reliability of military vehicles," *IEE Computing and Control Engineering*, 12:1, 11-20
- Neil, M., B. Malcolm, and R. Shaw, 2003, "Modelling an air traffic control environment using Bayesian belief networks," Presented at the 21st International System Safety Conference, August 4 - 8, Ottawa, Ontario, Canada
- Neil, M., N. Fenton, and M. Tailor, 2005, "Using Bayesian networks to model expected and unexpected operational losses," *Risk Analysis Journal*, 25, 963-972
- Neil, M., M. Tailor, and D. Marquez, 2007, "Inference in Bayesian networks using dynamic discretisation," *Statistics and Computing* 17:3
- Pearl, J., 1986, "Fusion, propagation, and structuring in belief networks," *Artificial Intelligence*, Vol. 29
- Rawnsley, J., 1995, *Going for broke: Nick Leeson and the collapse of Barings Bank*, HarperCollins, 1995.
- Reason, J., 1997, *Managing the risks of organisational accidents*, Ashgate Publishing Limited
- Spiegelhalter, D. J., and R.G. Cowell, 1992, "Learning in probabilistic expert systems," *Bayesian Statistics*, 4, 447-465
- The Times, 2008, "SocGen rogue trader 'ready to talk to police'," 25 January
- Venkataraman, S. S., 1997, "Value at risk for a mixture of normal distributions: the use of quasi-Bayesian estimation techniques," *Economic Perspectives*, Federal Reserve Bank of Chicago, 1997
- Wachtell, Lipton, Rosen, Katz & Promontory Financial Group, 2002, Report to the boards of Allied Irish Bank, p.l.c., Allfirst Financial Inc., and Allfirst Bank concerning currency trading losses March 12