

Ben Flood is a Research Fellow in the Centre for Telecommunications Value-Chain Research (CTVR) based in Trinity College Dublin, funded by Science Foundation Ireland (SFI). He received a (BSc) degree in Applied Mathematics and Computing from the University of Limerick (2002); and a Ph.D. in Statistics from the same institute (2008).

Ben Flood,
Department of Statistics,
Lloyd Building,
Trinity College Dublin,
Dublin 2,
Ireland.
floodbe@tcd.ie

Brett Houlding is a Research Fellow at CTVR based in Trinity College Dublin, funded by SFI. He received a (MMath) degree in Mathematical Sciences from Durham University (2004); and a Ph.D. in Statistics from the same institute (2008).

Brett Houlding,
Department of Statistics,
Lloyd Building,
Trinity College Dublin,
Dublin 2,
Ireland.
houldinb@tcd.ie

Simon Wilson is an Associate Professor in the Department of Statistics, Trinity College Dublin. He received a Ph.D. from George Washington University in 1993. He is a Fellow of the Royal Statistical Society and an elected member of the International Statistical Institute.

Simon P. Wilson,
Department of Statistics,
Lloyd Building,
Trinity College Dublin,
Dublin 2,
Ireland.
swilson@tcd.ie

Sergiy Vilkomir is a Research Associate Professor in the Software Quality Research Laboratory (SQRL), Department of Electrical Engineering and Computer Science, University of Tennessee. He received a Ph.D. in Computer Systems from the Kharkov Polytechnic Institute in 1990. He is a Senior Member of the IEEE.

Sergiy Vilkomir,
Software Quality Research Laboratory (SQRL),
Department of Electrical Engineering and Computer Science,
University of Tennessee,
Knoxville,
Tennessee,
USA.
vilkomir@cs.utk.edu

A Probability Model of System Downtime with Implications for Optimal Warranty Design

Ben Flood, Brett Houlding, Simon P. Wilson and Sergiy Vilkomir, Senior Member, IEEE.

Index Terms

System Availability, Bayesian Networks, Statistical Inference, Decision Theory, Segregated Failures Model, Warranty.

Abstract

Traditional approaches to modeling the availability of a system often do not formally take into account uncertainty over the parameter values of the model. Such models are then frequently criticised because the observed reliability of a system does not match that predicted by the model. Instead this paper extends a recently published segregated failures model so that, rather than providing a single figure for the availability of a system, uncertainty over model parameter values are incorporated and a predictive probability distribution is given. This predictive distribution is generated in a practical way by displaying the uncertainties and dependencies of the parameters of the model through a Bayesian network. Permitting uncertainty in the reliability model then allows the user to determine whether the predicted reliability was incorrect due to inherent variability in the system under study, or instead due to the use of an inappropriate model. Furthermore, it is demonstrated how the predictive distribution can be used when reliability predictions are employed within a formal decision-theoretic framework.

Use of the model is illustrated with the example of a high-availability computer system with multiple recovery procedures. A Bayesian network is produced to display the relations between parameters of the model in this case and to generate a predictive probability distribution of the system's availability. This predictive distribution is then used to make two decisions under uncertainty concerning offered warranty policies on the system: a qualitative decision, and an optimisation over a continuous decision space.

NOTATION

k	number of recovery levels
F	set of failure types
F_j	set of failures served at level j
F_{type_j}	set of failures of type j
p_{rec_j}	conditional probability of recovery at level j
p_{next_j}	conditional probability of progressing to service at level $j + 1$
p_{last_j}	conditional probability of immediately progressing to service at level k
p_{type_j}	probability that failure is of type j
τ_j	time taken to serve failure at level j (mins)
μ_j	restoration rate
λ	expected number of failures per year
λ_{type_j}	expected number of type j failures per year
v	actual number of failures in next year
v_{type_j}	actual number of type j failures in next year
τ_{type_j}	total downtime due to a type j failure (mins)
t_{d_j}	expected total yearly downtime due to failures of type j (mins)
T_d	expected total yearly downtime (mins)

This work was supported by grants from Science Foundation Ireland (SFI).

B. Flood (floodbe@tcd.ie), B. Houlding (houldinb@tcd.ie) and S.P. Wilson (swilson@tcd.ie) are with the Centre for Telecommunications Value-Chain Research (CTVR), Department of Statistics, Trinity College Dublin, Ireland.

S. Vilkomir is with the Software Quality Research Laboratory (SQRL), Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville, Tennessee, USA. He is a senior member of the IEEE.

c_j	coverage factor for recovery level j
$TN(\cdot, \cdot)$	Truncated Normal distribution
$\beta(\cdot, \cdot)$	Beta distribution
$u(\cdot)$	Utility function
$H(\cdot)$	Heaviside function
$E[\cdot]$	Expectation operator
$I_{\{\cdot\}}$	Indicator function

I. INTRODUCTION

THE quantification of uncertainty in mathematical models is essential to decision making. This paper combines data with expert knowledge in order to quantify the uncertainty over the input parameters of an availability model. A Bayesian network (BN) is then used to display the connections and dependencies between the parameters of the model and the combination of data and expert knowledge. The data and expert knowledge are combined in such a way that the uncertainty in the parameter values can be propagated through the mathematical availability model in order to generate a predictive distribution of the behaviour of the system. It is then shown how this prediction can in turn, with minimal effort, be incorporated into a formal decision-making framework.

The type of system to be considered here is such that, when system failure occurs, several recovery levels can be applied. It is assumed that some recovery procedures will be fast, but of limited application, whilst others will be slow, but applicable to a wider range of failures. In this setting a *recovery strategy* is considered to be a schedule of how the recovery procedures will be applied. A recovery strategy can be divided into several *levels*, each of which involves applying a separate recovery procedure.

In a fixed recovery strategy the order in which the recovery procedures will be applied is specified in advance. One choice of recovery strategy is to sequentially apply recovery levels so that the recovery level requiring the least amount of time is applied first. If following the application of this recovery level the system is not restored, then the next fastest recovery level is applied, and so on. It is presumed that the first few recovery levels will be applied automatically, and that the final recovery procedure, usually involving manual repair or replacement, is guaranteed to restore the system. However, for specific hardware failures it may be that automatic recovery levels are not likely to be successful. In this case all intermediate levels could be omitted, hence allowing the final recovery level to be immediately applied. An example of the use of such a recovery strategy is the Alcatel-Lucent Technologies Reliable Clustered Computing (RCC) product [1], which incorporates various recovery strategies in order to guarantee availability of commercial telecommunications systems.

There is a considerable literature on availability estimation in the situation when system parameters are known, but little for when their values are uncertain. However, quantifying this uncertainty is important if the model is to be used within a decision making analysis. As such this work addresses the particular issue of predicting the total downtime of a system, over the duration of a one year period, when the fixed recovery strategy is as described in [2] and [3], but where uncertainty is permitted over the correct value of the input parameters for the reliability model. In doing this, prior expert knowledge concerning likely values of input parameters is assumed available and used to model and quantify the uncertainty via relevant probability distributions. The effect that permitting uncertainty over parameter values has upon decision problems such as specifying a warranty price is also explored.

The number of parameters required in a suitable reliability model for a system can be large, and the connections and dependencies between their values can be complicated. In order to predict the availability of a system with variable behaviour, a joint probability distribution over all relevant parameters must be specified. However, identifying a suitable joint distribution immediately will often be a very difficult task. Instead a common technique in stochastic modeling is to determine the joint distribution through the use of a graphical model that displays the relationships between parameters (see, *e.g.*, [4]). A BN is a type of graphical model with the advantage that it facilitates the clear specification of the relationships

between system parameters and significantly reduces the computational cost of making statistical inference. Moreover, the versatility of BNs is becoming increasingly recognised in not only the field of reliability (see, for example, [5]–[9]), but also many other disciplines (*e.g.*, the modeling of genetics [10], and the modeling of financial systems [11]).

The remainder of this paper is structured as follows. Section II formally describes the segregated failure model and introduces an example system in which its use would be appropriate. Section III discusses how the example system may be modelled as a random process. Section IV introduces BNs, and shows how they are used when seeking to simulate the predictive distribution of the downtime in the example system. Section V discusses how data is incorporated and how the predictive posterior distribution is derived. Section VI discusses how the posterior predictive distribution is used within a decision making context concerning the selection of warranty costs. Finally, Section VII offers discussion and concluding remarks.

II. MODEL FOR SYSTEM RECOVERY

A. Model for System Failure

The segregated failure model is used in [2] and [3], and is as follows. Let there be k possible recovery levels, with recovery level k representing manual repair or replacement, and denote the set of all possible types of failure by $\mathbf{F}$. If a failure occurs in the system and recovery level j is applied, then the failure is said to have been *served* at level j . The set of failures that are served at level j is denoted as $\mathbf{F}_j \subseteq \mathbf{F}$. Unless the final recovery level is applied, recovery levels are applied consecutively in order of time required until successful restoration of the system has been achieved. Hence, for $j \neq k$, $f \in \mathbf{F}_j$ implies $f \in \mathbf{F}_l$, $0 < l < j$.

A failure f is said to be of type j if and only if recovery level j is the lowest level that will successfully recover the system from f . The set of all failures of type j is denoted as $\mathbf{F}_{type_j}$. Hence, $\mathbf{F}_{type_j} \subseteq \mathbf{F}_j$ and the sets $\mathbf{F}_{type_j}$ partition $\mathbf{F}$, *i.e.*,

$$\mathbf{F} = \bigcup_{j=1}^k \mathbf{F}_{type_j} \quad (1)$$

and

$$\mathbf{F}_{type_j} \cap \mathbf{F}_{type_l} = \emptyset, \forall j, l : 1 \leq j, l \leq k. \quad (2)$$

There are three possible outcomes if a failure is served at level j . The first possible outcome is that the failure is recovered. The conditional probability that a failure is recovered at level j given that it is served at level j will be denoted by p_{rec_j} , where

$$p_{rec_j} = \text{Prob}[f \in \mathbf{F}_{type_j} \mid f \in \mathbf{F}_j]. \quad (3)$$

The second possible outcome is that the system is not recovered and no specific hardware failure is identified. In this situation the failure is served at recovery level $j + 1$. The conditional probability of serving the failure at recovery level $j + 1$ given that it was served at recovery level j will be denoted by p_{next_j} , where

$$p_{next_j} = \text{Prob}[f \in \mathbf{F}_{j+1} \mid f \in \mathbf{F}_j]. \quad (4)$$

The final possible outcome is that the system is not recovered but that some specific hardware failure is identified, hence making manual repair or replacement necessary. In this situation the intermediate recovery levels $j + 1, \dots, k - 1$ are bypassed and the failure is served at recovery level k . The conditional probability that the failure is next served at recovery level k given that it has just been served at recovery level j will be denoted by p_{last_j} , where

$$p_{last_j} = 1 - p_{rec_j} - p_{next_j}. \quad (5)$$

Probability p_{rec_j} is defined for $1 \leq j \leq k$, whilst the description of p_{next_j} and p_{last_j} are only suitable for $1 \leq j \leq k-1$. It is assumed that the final recovery level always recovers the system, thus $p_{rec_k} = 1$. For consistency we fix $p_{last_k} = p_{next_k} = p_{last_{k-1}} = 0$ and $p_{next_{k-1}} = 1 - p_{rec_{k-1}}$.

There may also exist failures that can be escalated to the final recovery level before being served at recovery level 1. The probability that a failure will only be served at recovery level k will be denoted by p_{last_0} , whilst the probability that a failure will at least be served at recovery level 1 will be denoted by $p_{next_0} = 1 - p_{last_0}$.

Hence, the probability that a failure will be of type 1 is given by

$$p_{type_1} = p_{rec_1} p_{next_0}. \quad (6)$$

The probability that a failure will be of type j , $1 < j < k$, is given by

$$p_{type_j} = p_{rec_j} \left(\prod_{l=1}^{j-1} p_{next_l} \right) p_{next_0}. \quad (7)$$

The probability that a failure will only be recovered if it is served by the final recovery level k is given by

$$p_{type_k} = \left(p_{last_1} + \sum_{j=2}^{k-1} \left\{ p_{last_j} \prod_{l=1}^{j-1} p_{next_l} \right\} + \prod_{l=1}^{k-1} p_{next_l} \right) \times p_{next_0} + p_{last_0}, \quad (8)$$

or more simply

$$p_{type_k} = 1 - \sum_{j=1}^{k-1} p_{type_j}. \quad (9)$$

B. Model for System Downtime

It is assumed that the expected number of failures that occur in a given year is equal to the value of a parameter λ , and this will be referred to as the rate of failure. Also, the amount of time (measured in minutes) required to serve a failure at recovery level j will be denoted as τ_j . The reciprocal of τ_j , $\mu_j = 1/\tau_j$, is referred to as the restoration rate.

Fig. 1 below provides a graphical representation of the segregated failure model as given in [2].

In this setting every recovery level can be fully described by parameters $(\mu_j, p_{rec_j}, p_{next_j}, p_{last_j})$, and as such, a system with k recovery levels will be described by the $4k+2$ parameters $(\lambda, p_{next_0}, (\mu_j, p_{rec_j}, p_{next_j}, p_{last_j})_{j=1}^k)$.

The expected number of failures of type j that will occur in a given year is denoted as λ_{type_j} , and this is then found to be the following probability weighted proportion of the failure rate λ :

$$\lambda_{type_j} = \lambda p_{type_j}, \quad j = 1, \dots, k. \quad (10)$$

If there are v failures in a year, then we can decompose this number into the number of failures that were of type j , for $j = 1, \dots, k$. The number of failures that occur within a year that are of type j will be denoted as v_{type_j} and hence

$$v = \sum_{j=1}^k v_{type_j}. \quad (11)$$

The amount of time in minutes required to recover a failure of type j is denoted as τ_{type_j} , and as mentioned above, the service time τ_j is the time taken to serve a failure at recovery level j . Hence, the total recovery time in minutes that the system is down for due to a failure of type $j < k$ includes the




Fig. 1. The segregated failure model as given in [2].

time taken to serve the failure at recovery level j , τ_j , plus any additional time that was spent serving the failure at previous recovery levels, *i.e.*,

$$\tau_{type_j} = \sum_{l=1}^j \tau_l, j = 1, \dots, k-1. \quad (12)$$

For failures of type k the total recovery time will depend on whether or not the failure was served at recovery levels $j = 1, \dots, k-1$.

The expected overall downtime in minutes per year due to failures of type j will be denoted as t_{d_j} , and for $j < k$ is calculated as

$$t_{d_j} = \lambda_{type_j} \tau_{type_j}, j = 1, \dots, k-1. \quad (13)$$

The expected total downtime in minutes per year for the system, to be denoted as T_d , will then be

$$T_d = \sum_{j=1}^k t_{d_j}. \quad (14)$$

C. Example: Reliable Clustered Computing (RCC)

The RCC system is a collection of connected processors for which the segregated failure model is applicable as the recovery strategy. In particular, there are $k = 4$ possible recovery levels:

- 1) A switch-over from a failed active node to a spare node.
- 2) An automatic processor restart.
- 3) Reload data from disk, followed by an automatic processor restart.
- 4) Manual processor repair.

In [2] it is suggested that $p_{last_j} = 0$, for $1 \leq j < 4$. This means that a failure will only be diagnosed as being a serious hardware failure, hence requiring manual repair or replacement of the RCC, if it is diagnosed as such before being served at recovery level 1. However, a failure may indeed be of type 4

TABLE I
FAILURE PARAMETER VALUES

Param.	Value
λ	8 per year
τ_1	2 minutes
τ_2	5 minutes
τ_3	30 minutes
τ_4	240 minutes
c_a, c_1, c_2	0.9
c_3	0.99

without it being diagnosed as such, and if this happens the failure will be served at all recovery levels. Additional suggested parameter values for the segregated failure model are provided in Table I.

The ability of a recovery level to successfully restore a failure is called a *coverage factor*. The coverage factor for a recovery level j is defined to be the proportion of failures served at level j that are indeed fixed by recovery level j . In the RCC example the coverage factors c_1 , c_2 , and c_3 are the proportion of processor recoveries after switchover, processor restart, and restart with data reload, respectively. The value $1 - c_a$ is used to represent the proportion of serious hardware failures that are initially diagnosed as being such, and which are thus immediately served at recovery level 4.

The parameters p_{type_i} , $i = 1, \dots, 4$ used in the segregated failure model can be determined from the values of the coverage factors c_1 , c_2 , c_3 , and c_a in the following way:

$$p_{type_1} = c_1 c_a, \quad \text{from (6)} \quad (15)$$

$$p_{type_2} = c_2 (1 - c_1) c_a, \quad \text{from (7)} \quad (16)$$

$$p_{type_3} = c_3 (1 - c_1) (1 - c_2) c_a, \quad \text{from (7)} \quad (17)$$

$$p_{type_4} = c_a ((1 - c_1) (1 - c_2) (1 - c_3) - 1) + 1, \quad \text{from (8)} . \quad (18)$$

Suggested values for the coverage factors in this example are also available in Table I.

III. MODELING SYSTEM VARIATION: A STOCHASTIC MODEL

The downtime of most systems will inherently be variable. It may be true that for a given system the same operations will produce the same results every time, yet failures may still be governed by a random process such as, for example, the usage profile of the system. In practice the parameters of the system will not be known exactly, and a suitable and natural method for modeling such system variability is to permit uncertainty over the values of relevant parameters included in the failure model. Rather than specifying a fixed and presumed correct value for relevant parameters, uncertainty can formally be taken into account by instead eliciting probability distributions that are based upon expert knowledge. Such probability distributions concerning parameter values can then be used to generate a probability distribution concerning system availability.

One possible extension to the segregated failure model is to assume that failures can only occur when the system is not recovering, and that when the system is operational failures occur as though following a Poisson process with rate λ per year. Hence, assuming that the proportion of time the system is in recovery is relatively small, the number of yearly failures v follows a Poisson distribution with mean λ :

$$v | \lambda \sim \text{Poisson}(\lambda), \quad (19)$$

Here $v | \lambda$ is notation specifying that the stated distribution of v is conditional upon knowledge of the value of λ . Also note that the assumption of failures occurring as though following a Poisson process means that the duration times between failures are independent and that the expected number of failures over a given duration is constant, regardless of where in the year this duration is considered.

The rate that failures occur is specified by the parameter λ . However, it may be that there is uncertainty over λ 's value itself, in which case it too must be taken into account by specifying an appropriate probability distribution. To model the uncertainty over the rate of failure it is assumed that given parameters λ_m (affecting the expected failure rate) and λ_σ (affecting the level of uncertainty concerning the true value of the failure rate), λ follows a truncated Gaussian normal distribution, *i.e.*,

$$\lambda \mid \lambda_m, \lambda_\sigma \sim \text{TN}(\lambda_m, \lambda_\sigma^2). \quad (20)$$

Here the distribution is truncated at zero so as to prevent the possibility of a negative failure rate. Also note that although a Gaussian distribution is assumed for representing uncertainty over λ , it is trivial to replace this distribution with another that may better represent expert knowledge.

The segregated failure model also assumes that recovery levels $1, \dots, k-1$ (levels 1, 2, and 3 in the RCC example) are automated, and thus it may be reasonable to consider them as being stable processes with negligible variance in completion time. For this reason the values τ_j , $j = 1, \dots, k-1$, are assumed constant. The final recovery level, however, is likely to take a variable length of time to complete. Let τ_{k_i} be the i^{th} observed completion time for recovery level k . In order to model uncertainty over the completion time of recovery level k it is assumed that given parameters τ_{km} (affecting the mean completion time) and $\tau_{k\sigma}$ (affecting the amount of dispersion in completion times), τ_{k_i} also follows a truncated Gaussian normal distribution, *i.e.*,

$$\tau_{k_i} \mid \tau_{km}, \tau_{k\sigma} \sim \text{TN}(\tau_{km}, \tau_{k\sigma}^2). \quad (21)$$

This distribution is also truncated at zero so as to prevent the possibility of negative recovery times, and again it would be trivial to replace this distribution if another was believed to better represent expert knowledge.

Furthermore, we can permit uncertainty over the parameters τ_{km} and $\tau_{k\sigma}$, hence allowing observations of recovery times for recovery level k to affect beliefs over their values. The uncertainty over τ_{km} and $\tau_{k\sigma}$ will be assumed to be such that the use of a truncated Gaussian normal distribution is again appropriate, *i.e.*,

$$\tau_{km} \mid \tau_{kmm}, \tau_{km\sigma} \sim \text{TN}(\tau_{kmm}, \tau_{km\sigma}^2), \quad (22)$$

$$\tau_{k\sigma} \mid \tau_{k\sigma m}, \tau_{k\sigma\sigma} \sim \text{TN}(\tau_{k\sigma m}, \tau_{k\sigma\sigma}^2). \quad (23)$$

This is known as a hierarchical model (see, *e.g.*, [12]), where parameters of probability distributions are considered uncertain and are themselves assigned a probability distribution. The values τ_{kmm} , $\tau_{km\sigma}$, $\tau_{k\sigma m}$, and $\tau_{k\sigma\sigma}$ are assumed fixed and derived from expert knowledge. They affect the expected mean recovery time, the amount of uncertainty concerning mean recovery time, the expected amount of dispersion in recovery times, and the amount of uncertainty concerning dispersion in recovery times, respectively. Their values, however, will play an ever decreasing role in the model as data concerning observed final level recovery time is made available.

Finally, we can permit uncertainty over the coverage factors that are used to determine the probability that a failure is of a given type. Here it is assumed that the expert knowledge concerning the coverage factors in the RCC example is such that a Beta probability distribution is appropriate with specified and fixed prior parameters, *i.e.*,

$$c_i \mid c_{i\alpha}, c_{i\beta} \sim \beta(c_{i\alpha}, c_{i\beta}), i = 1, 2, 3, \quad (24)$$

$$c_a \mid c_{a\alpha}, c_{a\beta} \sim \beta(c_{a\alpha}, c_{a\beta}). \quad (25)$$

As with all the distributions discussed in this section, it would be trivial to replace these if they were not in agreement with expert knowledge.

IV. BAYESIAN NETWORKS

A Bayesian network (BN) is a directed acyclic graph and is used to provide a graphical representation of the dependencies between the set of random variables within a probability model (see, *e.g.*, [13]). The network represents causal relationships between the variables and there is one node for each variable in the problem. A directed edge from node A to node B represents the situation where the variable at node A ‘causes’ what happens to the variable at node B . More formally, a directed edge from A to B represents the situation in which the probability distribution of the variable at B is a function of the value of the variable at node A .

A key property of a BN is that the joint probability distribution of all the variables in a network can be decomposed from it. If the network contains n random variables $X_1, \dots, X_n$, then their joint probability distribution can be simplified to

$$p(X_1 = x_1, \dots, X_n = x_n) = \prod_{i=1}^n p(X_i = x_i | \text{parents of } X_i). \quad (26)$$

A parent of X_i is a variable whose node has an edge that leads to the X_i node.

Fig. 2 shows the Bayesian network of the subset of the RCC system corresponding to only the first recovery level. Rectangular nodes represent parameters with fixed, known, values. In practice, however, many parameters will have an uncertain value, and these are represented in the network by circular nodes. Following the usual Bayesian approach, prior distributions should be specified for such uncertain parameters by referring to expert knowledge of the system.




Fig. 2. The BN describing the variability of the subset of the system corresponding to failures of type 1. The BN includes uncertainty about λ , but assumes perfect knowledge about p_{type_1} and τ_{type_1} . The prior parameters for λ are constant and are thus represented by rectangular nodes.

Fig. 2 illustrates that the total downtime of the system due to failures of type 1, t_{d_1} , is uncertain, but has a probability distribution that depends only on the known duration of time required to fix a failure of type 1, τ_{type_1} , and the number of type 1 failures that occur over the year, v_{type_1} (provided that these values

are known). However, v_{type_1} is itself uncertain, but has a probability distribution that depends only on the known probability that a failure will be of type 1, p_{type_1} , and the total number of failures that occur over the year v (again provided that these values are known). The total number of failures is, however, also uncertain, but has a probability distribution that depends only on the failure rate λ when that parameter is known. Finally, λ is also uncertain, but has a probability distribution depending only on the known prior parameters λ_m and λ_σ .

The BN can hence be used to decompose the joint probability distribution in the following way:

$$\begin{aligned} p(t_{d_1}, v_{type_1}, v, \lambda \mid \tau_{type_1}, p_{type_1}, \lambda_m, \lambda_\sigma) \\ = p(t_{d_1} \mid v_{type_1}, \tau_{type_1}) p(v_{type_1} \mid v, p_{type_1}) p(v \mid \lambda) p(\lambda \mid \lambda_m, \lambda_\sigma). \end{aligned} \quad (27)$$

With the exception of $p(v_{type_1} \mid v, p_{type_1})$, all of the distributions on the right hand side of Equation (27) have already been specified. In order to model the uncertainty concerning v_{type_j} , $j = 1, \dots, k$, it is assumed that $(v_{type_1}, \dots, v_{type_k})$ follows a multinomial distribution with parameters $(p_{type_1}, \dots, p_{type_k})$ and sample size v .

V. INCORPORATING DATA

A. Including Data in the BN

As data from the observation of failures in previous years or identical systems are observed, the uncertainty concerning the system's parameter values would be expected to decrease, and Bayes' theorem provides us with a convenient formula for updating beliefs concerning the system's parameters. If v is regarded as the uncertain number of system failures that will occur in the next year, then knowledge of $(v_1, v_2, \dots, v_{n_v})$, with v_i the number of failures previously observed over the duration of a year for an identical system i , will clearly be relevant in determining beliefs over v . If such data is available, then using the structure of the BN and Equation (26) the joint distribution of $(v_1, v_2, \dots, v_{n_v})$, v and λ can be determined as

$$p(v_1, v_2, \dots, v_{n_v}, v, \lambda \mid \lambda_m, \lambda_\sigma) = p(v \mid \lambda) p(\lambda \mid \lambda_m, \lambda_\sigma) \prod_{i=1}^{n_v} p(v_i \mid \lambda). \quad (28)$$

Bayes' theorem then informs us that the probability distribution of v given knowledge of the data $v_1, \dots, v_{n_v}$ is proportional to $\int p(v \mid \lambda) p(\lambda \mid \lambda_m, \lambda_\sigma) \prod_{i=1}^{n_v} p(v_i \mid \lambda) d\lambda$.

The BN can thus be extended to have one node for each data observation v_i , $i = 1, \dots, n_v$. However, a tidier way of representing these variables within the BN is to use a shaded rectangular node (rectangular as the data is known and not uncertain), with the shading symbolising that more than one observation of the same system attribute is represented. Fig. 3 shows the subgraph of the BN that takes into account the model of failure occurrence.

The process of generating the posterior distribution of v given the data $(v_1, v_2, \dots, v_{n_v})$ is easily seen through examination of the BN. First samples must be drawn from the posterior distribution of λ given the data $(v_1, v_2, \dots, v_{n_v})$. After this samples may be drawn from distribution $p(v \mid \lambda)$. Depending on the form of the probability distributions employed within the model, the posterior distribution may be found in closed form. However, alternative methods for determining the posterior distribution exist if this is not possible. Here the adaptive rejection Metropolis sampling (ARMS) technique has been used to simulate from all posterior distributions. This technique was implemented using the `arms` [14] software package within the software programme R [15].

If along with the data concerning previous failures it was noted that there had been n_{τ_k} failures served at the final recovery level, and that the duration of time required for each of these was recorded, then a posterior distribution may also be found in a similar way for the time in which it takes to apply the final recovery level τ_k .




Fig. 3. BN including n_v observations of annual number of failures, $(v_1, v_2, \dots, v_{n_v})$, and prior uncertainty for λ . The data is represented by a rectangular shaded node to make the graph more tidy.

B. A Complete BN for the Reliable Clustered Computing Model

The full BN for the segregated failure model as applied to the RCC system is now presented in Fig. 4. This diagram facilitates easy identification of the joint distribution of the variables involved. Furthermore, once the joint distribution has been identified, common Bayesian statistical methods can be applied in order to determine the posterior predictive distribution of any variables of interest, and in particular for this model, total annual downtime.

The BN of Fig. 4 demonstrates that the RCC failure model has a joint probability distribution as specified in Equation (29). The right hand side of Equation (29) consists of a product of distributions, each of which has been specified within the model. The proportionality sign is included because of the omission of known distributions that are actually deterministic relations, e.g., $p(p_{type1} | c_1, c_a)$.

Let $\mathcal{P} = \{c_a, c_1, c_2, c_3, p_{type1}, \dots, p_{type4}, \lambda, v, v_{type1}, \dots, v_{type4}, \tau_{4m}, \tau_{4\sigma}, \tau_{type41}, \dots, \tau_{type4v_{type4}}, t_{d1}, t_{d2}, t_{d3}, t_{d4}, T_d\}$ be the set of uncertain parameter values in the BN of Fig. 4. In addition, let $\mathcal{C} = \{\lambda_m, \lambda_\sigma, c_{a\alpha}, c_{1\alpha}, c_{2\alpha}, c_{3\alpha}, c_{a\beta}, c_{1\beta}, c_{2\beta}, \dots\}$ be the set of fixed and constant values and let $\mathcal{D} = \{v_i\}_{i=1}^{n_v} \cup \{\tau_{4i}\}_{i=1}^{n_{\tau_4}}$ be the observed data. Then,

$$\begin{aligned}
 p(\mathcal{P}, \mathcal{D} | \mathcal{C}) &\propto p(\lambda | \lambda_m, \lambda_\sigma) p(v | \lambda) \left(\prod_{i=1}^{n_v} p(v_i | \lambda) \right) p(c_a | c_{a\alpha}, c_{a\beta}) \left(\prod_{i=1}^3 p(c_i | c_{i\alpha}, c_{i\beta}) \right) \\
 &\quad \times p(v_{type1}, \dots, v_{type4} | v, p_{type1}, \dots, p_{type4}) p(\tau_{4m} | \tau_{4m_m}, \tau_{4m_\sigma}) p(\tau_{4\sigma} | \tau_{4\sigma_m}, \tau_{4\sigma_\sigma}) \\
 &\quad \times \left(\prod_{i=1}^{n_{\tau_4}} p(\tau_{4i} | \tau_{4m}, \tau_{4\sigma}) \right) p(\tau_{type41}, \dots, \tau_{type4v_{type4}} | v_{type4}, \tau_{4m}, \tau_{4\sigma}, c_a, \tau_{type3}).
 \end{aligned} \tag{29}$$

As mentioned previously, $p(v_{type1}, \dots, v_{type4} | v, p_{type1}, \dots, p_{type4})$ follows a multinomial distribution, and the expression $p(\tau_{type41}, \dots, \tau_{type4v_{type4}} | v_{type4}, \tau_{4m}, \tau_{4\sigma}, c_a, \tau_{type3})$ can be simplified further when it is realised that each τ_{type4i} is an independent sample that with probability $1 - c_a$ is from the distribution $TN(\tau_{4m}, \tau_{4\sigma})$, or otherwise is from the distribution $TN(\tau_{4m}, \tau_{4\sigma}) + \tau_{type3}$.

Equation (29) looks rather complicated, and indeed would be more so if instead the full expression that includes the deterministic relations had been given. However, and as discussed previously, once the joint distribution has been identified, its use allows easy determination of the posterior distribution for any parameters of interest.

The procedure used here for determining posterior distributions involves sampling from the joint distribution of the parameters and data. To do this the variables whose distribution only depends on




Fig. 4. The full BN for modeling total yearly downtime of the RCC system. The simulation method to determine the predictive distribution is relatively straightforward to identify once this graph is defined. Notice that the graph does not specify the distributions of individual variables, making it appropriate even if expert knowledge is not represented by the distributions assumed here.

known constants are sampled first. Following this samples are drawn of the variables whose distribution only depends on known constants and those variables that have already been sampled, and so on. Taking enough samples then allows an approximation of any posterior distribution of interest. The only difficulty in the case of the RCC example and the BN of Fig. 4 is that τ_{4m} and $\tau_{4\sigma}$ are both linked to the constant data values $\tau_{41}, \tau_{42}, \dots, \tau_{4n_{\tau_4}}$, hence leading to dependency between the samples of τ_{4m} and $\tau_{4\sigma}$. This results in an inseparable posterior distribution $p(\tau_{4m}, \tau_{4\sigma} \mid \tau_{41}, \tau_{42}, \dots, \tau_{4n_{\tau_4}})$. Nevertheless, posterior samples of these variables may still be determined through the use of Gibbs sampling (see, *e.g.*, [12]), whereby samples are sequentially taken in turn from the conditional distributions $p(\tau_{4m} \mid \tau_{4\sigma}, \tau_{41}, \tau_{42}, \dots, \tau_{4n_{\tau_4}})$ and $p(\tau_{4\sigma} \mid \tau_{4m}, \tau_{41}, \tau_{42}, \dots, \tau_{4n_{\tau_4}})$.

As mentioned previously, in demonstrating the concept of permitting uncertainty over parameters and determining the joint distribution through the use of a BN, it was assumed that expert knowledge was such that the prior distributions for λ , τ_{4m} and $\tau_{4\sigma}$ were all truncated Normal distributions. The priors for the coverage factors c_a , c_1 , c_2 , and c_3 were all assumed to be beta distributions and the variables $v_{type_1}, \dots, v_{type_4}$ are assumed to be multinomially distributed given the parameters $p_{type_1}, \dots, p_{type_4}$ and sample size v . The values of the other parameters within the BN can be determined functionally from these.

The required initial constants for the BN are assumed to be as given in Table II, with these values being derived from the values used in [2].

TABLE II
PRIOR PARAMETERS

Parameter	Value	Parameter	Value
λ_m	8	c_{a_α}	9
λ_σ	2	c_{a_β}	1
τ_{4m_m}	240	c_{1_α}	9
τ_{4m_σ}	20	c_{1_β}	1
$\tau_{4\sigma_m}$	20	c_{2_α}	9
$\tau_{4\sigma_\sigma}$	10	c_{2_β}	1
τ_1	2 mins	c_{3_α}	99
τ_2	5 mins	c_{3_β}	1
τ_3	30 mins		

In order to determine posterior distributions data concerning previous failures is required. In place of this, however, it is assumed that the data concerning the number of previous annual failures for 10 identical systems equals 10 samples from a Poisson distribution with mean $\lambda = 8$. It is also assumed that observations of repair times for the final recovery level equal samples from the truncated Normal distribution $TN(240, 8^2)$. An example of the type of data used to determine posterior distributions is given in Table III.

TABLE III
RELIABILITY DATA FOR 10 IDENTICAL SYSTEMS

Number of Failures	Final Level Repair Times (mins)
8	233, 223
10	251
9	242
8	239
4	236, 247
10	250, 242
9	-
7	241, 247, 249
11	244
3	233, 243

An approximation of the posterior predictive distribution for total annual downtime T_d based on 10,000 samples is provided in Fig. 5. Both the density function in Fig. 5 and the density function in Fig. 7 are estimated with the use of Gaussian kernels.

The multi-modal shape of the probability density function of Fig. 5 is caused by the length of time required for manual repair, with each peak corresponding to the number of manual repairs required in a year. The predictive distribution of T_d is also highly skewed, with the mean value of T_d being 198 minutes and the median value being 20 minutes. This corresponds to a mean predicted proportion of availability time of 0.9996252 and a median predicted proportion of availability time of 0.999962.

Further features of the system can also be established using the simulation of the joint posterior distribution, and Table IV offers some summary statistics.

Additional results of the sampling procedure show that there were some samples in which no failure of the system occurred, hence leading to a minimum annual downtime of zero. Moreover, the occurrence of a zero median value for the total downtime due to a given failure type in Table IV means that in more than half of the 10,000 simulations there was no observation of that failure type, and this is found to be the case for failure types 2, 3 and 4.

The joint posterior distribution can also be used to answer additional questions of interest. For example, it might be important to determine which failure types are causing the most extreme losses of availability,




Fig. 5. The posterior predictive distribution for downtime of the system over one year given data as in Table III. Units for the x -axis correspond to minutes per year.

TABLE IV
POSTERIOR DISTRIBUTION OF MINUTES/YEAR FOR EACH TYPE OF DOWNTIME.

	2.5% Percentile	Median	Mean	97.5% Percentile
t_{d1}	2	12	12.85	24
t_{d2}	0	0	4.473	21
t_{d3}	0	0	0.977	14
t_{d4}	0	0	197.3	934

and hence it may be useful to know the distribution of failures when the total annual downtime is greater than 300 minutes. This distribution can be found by restricting attention to only those samples in which it is found that $T_d > 300$, and summary statistics for this situation are provided in Table V.

TABLE V
DISTRIBUTION OF MINUTES/YEAR FOR EACH TYPE OF DOWNTIME $T_d > 300$.

	2.5% Percentile	Median	Mean	97.5% Percentile
t_{d1}	2	12	12.24	24
t_{d2}	0	0	4.534	21
t_{d3}	0	0	1.027	14
t_{d4}	244.2	528	615	1325

Both Tables IV and V demonstrate that the main cause of downtime is the requirement for manual repair.

Finally, rather than examining only the estimate for the expected annual downtime of the system, an idea of the level of variation in annual downtime is also important. To determine prediction intervals for T_d the simulated values are placed in increasing order and percentiles noted. For example, the fifth and ninety-fifth percentile of the simulated sample provides a 90% prediction interval, which in this case is found to be [8mins, 804mins]. Such statistics over variation in downtime would not be available if uncertainty over the model's parameters had not been explicitly included.

VI. DECISION MAKING UNDER UNCERTAINTY

A. Employing the Posterior Distribution within a Decision Analysis

Obtaining the joint posterior probability distribution of all the variables in the failure model is not necessarily the end goal. Rather, the objective will often be to use such information within a decision analysis, *e.g.*, in deciding the warranty cost for the system. To achieve this the problem is framed as one of decision making under uncertainty, the theory of which falls neatly in line with that of Bayesian statistical analysis (see, *e.g.*, [16]). However, in order to specify the decision problem a utility function is required that describes the worth of each possible outcome. Once this has been identified, the theory stipulates that the decision maximising expected utility should be selected. The outcome of the decision problem in this instance is assumed to be monetary return.

For a set of possible outcomes $\mathcal{X}$, a utility function is defined as a map $u : \mathcal{X} \rightarrow \mathbb{R}$ that is in agreement with the decision maker's preferences, *i.e.*, if $X_1, X_2 \in \mathcal{X}$ are two possible outcomes with the property that $u(X_1) \geq u(X_2)$, then the decision maker holds that outcome X_1 is at least as preferable as outcome X_2 .

As an example of the type of decision analysis that may be relevant consider the situation in which the supplier of the system offers a warranty policy that includes stated penalties to be paid depending on the system's performance over its first year of use. In particular, suppose that if the annual downtime is greater than 600 minutes, a penalty of $\$200 + (T_d - 600) \times \10 is paid, whilst if the annual downtime is between 400 and 600 minutes, a penalty of $\$200$ is paid. However, no penalty is to be paid if the annual downtime is less than 400 minutes.

Presuming that the appropriate utility function for outcomes that take the form of monetary returns is a linear function, the utility function for this warranty policy is a function of total annual downtime and is such that

$$u(T_d) = -200 (H(T_d - 400) - H(T_d - 600)) - H(T_d - 600) (200 + 10(T_d - 600)). \quad (30)$$

Here $H(\cdot)$ represents the usual Heaviside function. Also note that this is the utility function for the supplier of the system. Fig. 6 plots the financial return to the supplier for various values of T_d .




Fig. 6. The warranty policy depends on the observed downtime over the first year of use of the system with the function $u(T_d) = -200(H(T_d - 400) - H(T_d - 600)) - H(T_d - 600)(200 + 10(T_d - 600))$ representing both cost and utility due to the assumption of linearity.

The approximated posterior predictive distribution of annual downtime T_d can now be used to determine the expected financial return to the supplier of the warranty policy after one year, and as here it is assumed

that utility is a linear function of monetary returns, can be used to determine the expected utility for offering such a warranty on this system. If L samples $T_{d_1}, \dots, T_{d_L}$ are taken of the annual downtime T_d , then the expected cost can be approximated as

$$\frac{\sum_{i=1}^L u(T_{d_i})}{L}. \quad (31)$$

However, suppose the supplier of the system has the option to sub-contract all manual repairs at a fixed cost of \$100 per repair, but where each such repair is then guaranteed to take precisely 150 minutes. In this situation the supplier faces the problem of deciding whether or not to agree to such a sub-contract. Let O_1 refer to the option of not sub-contracting and let O_2 refer to the option of sub-contracting.

A BN for the model of system failure under O_2 can again be provided, and in conjunction with the stated utility function, be used to determine whether or not the supplier should sub-contract manual repair. Indeed, the appropriate BN for O_2 is a simplification of that given in Fig. 4, and corresponds to reducing the set of parents of $\tau_{type_{4_i}}$ to $\{v_{type_{4_i}}, c_a, \tau_{type_{3_i}}\}$.

A simulated posterior predictive probability distribution function for total annual downtime T_d under O_2 is provided in Fig. 7. The multi-modal structure of this distribution is similar to that of the distribution under O_1 , but now the modes are closer together and there is less probability mass between them.




Fig. 7. Predicted distribution of annual downtime of the system with fixed manual repair time. Units for the x -axis correspond to minutes per year.

Taking a simulation of $L = 10,000$ samples of T_d under O_1 , $T_{d_1}, \dots, T_{d_{L_1}}$, the expected utility of this option is approximated to be

$$u(O_1) = \sum_{i=1}^L \frac{u(T_{d_{i_1}})}{L} = -\$329. \quad (32)$$

To determine the expected utility of option O_2 , the fixed manual repair costs must also be taken into account. Taking a simulation of $L = 10,000$ samples of T_d under O_2 , $T_{d_{1_2}}, \dots, T_{d_{L_2}}$, the expected utility of this option is approximated to be

$$u(O_2) = \sum_{i=1}^L \frac{u(T_{d_{i_2}}) - 100v_{type_{4_i}}}{L} = -\$40. \quad (33)$$

Hence fixing the manual repair time to be 150 minutes at an additional cost of \$100 per repair of this type reduces the expected cost of the warranty policy by \$289.

B. Continuous Decision Space

Another interesting challenge is to alter aspects of the warranty policy so as to maximise expected utility/financial return. Maintaining the same overall shape of the original warranty policy, the utility function can be parameterised using the parameters $\phi_1, \dots, \phi_4$ so that

$$u(T_d, v_{type_4}, \phi_1, \dots, \phi_4) = -\phi_3(H(T_d - \phi_1) - H(T_d - \phi_2)) - H(T_d - \phi_2)(\phi_3 + \phi_4(T_d - \phi_2)) - 100v_{type_4}. \quad (34)$$

Here ϕ_1 is the maximum amount of downtime permitted over the year before a penalty is incurred, ϕ_2 is the maximum amount of annual downtime that is permitted for a single fixed penalty to be incurred, ϕ_3 is the fixed penalty incurred for when $\phi_1 \leq T_d \leq \phi_2$, and ϕ_4 is the additional penalty incurred for each minute of annual downtime beyond $T_d = \phi_2$. These parameters are restricted to be such that $\phi_i \in [0, \infty)$, for $i = 1, \dots, 4$, and such that $\phi_2 \geq \phi_1$.

Of course as the problem stands, maximum utility is guaranteed when $\phi_3 = \phi_4 = 0$. However, providing such a warranty policy would clearly affect sales of the system, as in effect it states that no penalty will be incurred by the system supplier regardless of the system's total annual downtime. Instead, and to make the problem more realistic, it is assumed that the values of $\phi_1, \dots, \phi_4$ affect the likelihood that the system is purchased by the customer.

Let B be the event that a customer buys the system and B^c the complement event where the customer does not buy the system. If B occurs, then the product is sold for a profit P , but will also be subject to an uncertain warranty cost $u(T_d, v_{type_4}, \phi_1, \dots, \phi_4)$, whilst if B^c occurs there is no profit or loss. The expected utility return for a given selection of warranty parameters $\phi_1, \dots, \phi_4$ is thus:

$$U(\phi_1, \dots, \phi_4) = P(B|\phi_1, \dots, \phi_4) (P + u(T_d, v_{type_4}, \phi_1, \dots, \phi_4)) \quad (35)$$

The decision problem is thus to select values of $\phi_1, \dots, \phi_4$ so as to maximise the expectation of $U(\phi_1, \dots, \phi_4)$.

$$\begin{aligned} E[U(\phi_1, \dots, \phi_4)] &= E[p(B|\phi_1, \dots, \phi_4) (P + u(T_d, v_{type_4}, \phi_1, \dots, \phi_4))] \\ &= p(B|\phi_1, \dots, \phi_4) \left(P + \int u(T_d, v_{type_4}, \phi_1, \dots, \phi_4) p(T_d, v_{type_4}) dT_d dv_{type_4} \right). \end{aligned} \quad (36)$$

The expression $Int = \int u(T_d, v_{type_4}, \phi_1, \dots, \phi_4) p(T_d, v_{type_4}) dT_d dv_{type_4}$ can be further simplified as

$$Int = -\phi_3 \times p[T_d > \phi_1] - \phi_4 E[T_d - \phi_2 | T_d > \phi_2] - 100E[v_{type_4}]. \quad (37)$$

To proceed a form for the probability $p(B|\phi_1, \dots, \phi_4)$ is required, and whilst this is a difficult elicitation problem in itself, here it is assumed that this distribution takes the logistical shape (where $I_{\{x\}}$ is the indicator function that takes value 1 when x is true and value 0 otherwise)

$$p(B|\phi_1, \dots, \phi_4) = \frac{1}{1 + e^{-\delta - \delta_1 \phi_1 I_{\{\phi_3 > 0\}} - \delta_2 \phi_2 I_{\{\phi_4 > 0\}} - \delta_3 \phi_3 - \delta_4 \phi_4}}. \quad (38)$$

Data from customer surveys, past sales experience, or expert knowledge can be used to determine the values of the fixed constants $\delta, \delta_1, \dots, \delta_4$.

Hence, the problem is to determine

$$\arg \max_{\phi_1, \dots, \phi_4} \left\{ \frac{P - \phi_3 \times p[T_d > \phi_1] - \phi_4 E[T_d - \phi_2 | T_d > \phi_2] - 100E[v_{type_4}]}{1 + e^{-\delta - \delta_1 \phi_1 - \delta_2 \phi_2 - \delta_3 \phi_3 - \delta_4 \phi_4}} \right\} \quad (39)$$

To solve this problem the samples from the BN can again be made use of and, given values for the constants P , δ , and $\delta_1, \dots, \delta_4$, optimal values for the parameters $\phi_1, \dots, \phi_4$ can be found easily by using the `optim` function in R.

For example, assuming $(P, \delta, \delta_1, \delta_2, \delta_3, \delta_4) = (1000, 2, -0.003, -0.01, 0.000008, 2)$ the optimal parameter values are approximated to be $(\phi_1, \phi_2, \phi_3, \phi_4) = (194.0, 194.9, 252.0, 1.2)$, leading to an approximated maximum expected return of \$543. Of note in this instance is that the feature of allowing a range of system downtime in which only a fixed cost is incurred provides little additional benefit, with the range being only 0.9 minutes in the optimal policy.

The important distinction, however, is that without incorporating uncertainty into the model, only an estimate of expected downtime would be available, rather than the predictive distribution that is required in a formal decision-theoretic setup. The expected downtime for a system can have zero penalty under a specific warranty policy, yet this is unlikely to be the true expected cost of offering that particular policy. Indeed, given the values employed in the above example it is found that expected downtime is 159.5 minutes, below the 194 minutes of downtime that must occur before any penalty is incurred.

Only if uncertainty over the value of annual system downtime is taken into account can the true expected warranty cost be determined, a clear motivation for the use of this modeling approach, especially if it is deemed more important to have knowledge of the true expected warranty cost than to simply have knowledge of the expected annual downtime.

VII. DISCUSSION AND CONCLUSIONS

When mathematically modeling a system it is common practice to return a constant output parameter for different values of the input parameters. Instead in this paper the use of prior distributions concerning expert knowledge has been utilised to fully specify the knowledge available about the state of each parameter before calculating the model. It was then shown how this knowledge can be updated using available data.

The approach of representing dependencies through a BN permitted easy determination of necessary predictive distributions by sampling from the relevant joint distribution of the system parameters. For more complex systems, however, the speed of such a simulation may become important. If this is indeed the case then it may be beneficial to select prior distributions that, whilst not being a perfect representation of expert knowledge, allow easy determination of posterior distributions. Further examination of the appropriate utility function could also be included, for whilst here it is assumed that a linear function was representative of true preferences, there is no reason why this can not be changed to another if it were deemed more reasonable, *e.g.*, a logarithmic function so as to capture feelings of risk-aversion.

The clear difficulty to the approach outlined here is that it is heavily dependent on information from domain experts, whose time will undoubtedly be of value. Nevertheless, if such information is available, then it should be included in order to generate a model that fits closely with observed results.

Finally, possible extensions are to consider the use of continuous-time Bayesian networks (CTBNs) for modeling system dependencies (see, *e.g.*, [5]), and to incorporate a cost for data collection that could then be included in the formal decision problem. A possibility for the latter of these would be the method of sampling the decision space as used in, for example, [17].

REFERENCES

- [1] G. Hughes-Fenchel, "A flexible clustered approach to high availability," *27th Annual International Symposium on Fault-Tolerant Computing, FTCS-27, Seattle, Washington*, pp. 314–318, Jun. 1997.
- [2] S. A. Vilkomir, D. L. Parnas, V. B. Mendiratta, and E. Murphy, "Availability evaluation of hardware/software systems with several recovery procedures," *Proc. 29th IEEE Annual International Computer Software and Applications Conference, Edinburgh, Scotland*, pp. 473–478, Jul. 2005.
- [3] —, "Segregated failures model for availability evaluation of fault-tolerant systems," in *Proc. 29th Australasian Computer Science Conference (ACSC 2006)*, vol. 48, Tasmania, Hobart, Australia, Jan. 2006, pp. 55–61.
- [4] L. Lauritzen, Steffen, *Graphical Models*, 2nd ed., ser. Oxford Statistical Science Series. New York: Oxford University Press Inc., 1998.

- [5] H. Boudali and J. Dugan, "A continuous-time bayesian network reliability modeling, and analysis framework," *IEEE Transactions on Reliability*, vol. 55, no. 1, pp. 86–97, Mar. 2006.
- [6] A. Helminen, "Reliability estimation of safety-critical software-based systems using bayesian networks," Radiation and nuclear safety authority of Finland (STUK), Tech. Rep., 2001. [Online]. Available: <http://www.stuk.fi/julkaisut/tr/stuk-yto-tr178.pdf>
- [7] M. Hiirsalmi, "Method feasibility study: Bayesian networks," VTT Information Technology, Tech. Rep., 2000. [Online]. Available: <http://virtual.vtt.fi/inf/julkaisut/muut/2000/rr2k29-c-ww-feasib.pdf>
- [8] B. Littlewood, L. Strigini, and D. Wright, "Assessing the reliability of diverse fault-tolerant systems," in *Proc. INucE International Conference on Control and Instrumentation in Nuclear Installations*, Bristol, UK, 2000.
- [9] —, "Examination of bayesian belief network for safety assessment of nuclear computer-based systems," DeVa, City University, London, Tech. Rep., 1998. [Online]. Available: http://www.csr.city.ac.uk/people/lorenzo.strigini/ls.papers/DeVa_BBN_reports/DeVaTR70_year3.5a/DeVaTR70.pdf
- [10] S. L. Lauritzen and N. A. Sheehan, "Graphical models for genetic analyses," *Statistical Science*, vol. 18, no. 4, pp. 489–514, 2003.
- [11] J. Gemela, "Financial analysis using bayesian networks," *Applied Stochastic Models in Business and Industry*, vol. 17, pp. 57–67, 2001.
- [12] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis, Second Edition*. Chapman & Hall/CRC, 2003.
- [13] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988.
- [14] G. Petris, L. Tardella, and W. Gilks., *HI: Simulation from distributions supported by nested hyperplanes: R package version 0.1-1*, 2005.
- [15] R Development Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2006, ISBN 3-900051-07-0. [Online]. Available: <http://www.R-project.org>
- [16] D. V. Lindley, "The philosophy of statistics," *Journal of the Royal Statistical Society, Series D: The Statistician*, vol. 49, no. 3, pp. 293–319, 2000.
- [17] P. Müller, B. Sansó, and M. De Iorio, "Optimal Bayesian design by inhomogeneous Markov chain simulation," *Journal of the American Statistical Association*, vol. 99, no. 467, pp. 788–798, 2004.