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 has submitted a Ph.D. in Statistics to the same institute (2006).

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

Simon Wilson is a senior lecturer in the Department of Statistics, Trinity College Dublin. He received a Ph.D. in Stochastic Modeling from the 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.

Sergiy Vilkomir is a senior researcher in Software Quality Research Laboratory (SQRL), Department of Computer Science and Information Systems, University of Limerick. He received a Ph.D. in Computer Systems from the Kharkov Polytechnic Institute in 1990.

Sergiy Vilkomir,
Software Quality Research Laboratory (SQRL),
Department of Computer Science and Information Systems,
Faculty of Informatics and Electronics,
University of Limerick, Limerick,
Ireland.

Propagation of Uncertainty Through a Segregated Failure Model

Ben Flood, Simon P. Wilson, and Sergiy Vilkomir

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

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

S. Vilkomir is with the 1Software Quality Research Laboratory (SQRL), Department of Computer Science and Information Systems, University of Limerick, Limerick, Ireland.

Propagation of Uncertainty Through a Segregated Failure Model

Index Terms

Bayesian Networks, Decision Theory, Segregated Failures Model.

I. Summary and Conclusions

This paper takes a recently published segregated failures model that gives a single figure for the availability for a system, and generates a predictive distribution for the availability for the system. The predictive distribution is generated in a practical way by modelling the system using a Bayesian network. The predictive distribution is then used to make two decisions, the first a qualitative decision, and the second an optimisation over a continuous decision space.

Reliability models are often criticised because the observed reliability of a system does not match the predicted reliability. This work is useful because it allows the generation of predictive distributions, which model the uncertainty in the prediction of reliability. This can be used to determine if the predicted reliability was incorrect due to the use of an incorrect model or due to variability inherent to the system under study. Further to this, the predictive distribution can be used to place reliability predictions under uncertainty into a formal decision-theoretic framework.

II. Introduction

Quantification of uncertainty in mathematical models is essential to decision making. This paper uses a Bayesian Network approach to combine data with expert knowledge to quantify the uncertainty of the input parameters of an availability model. 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 to generate a predictive distribution of the behaviour of the system. It is shown how this prediction can in turn, with minimal effort, be incorporated into a decision-making framework.

The system under study here is designed such that, when it fails, several recovery levels can be applied. Some recovery procedures will be fast, but of limited application. Others may be

slow, but applicable to a wider range of failures. A *recovery strategy* is a schedule of how the recovery procedures will be applied. A recovery strategy can be divided into several *levels*, each of which involves 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 apply the recovery level that takes the shortest amount of time first. If the system is not restored in that recovery level, then the next fastest recovery level is applied. The first few recovery levels can usually be applied automatically. The final recovery procedure, usually involving manual repair or replacement, is guaranteed to restore the system. For specific hardware failures it may be diagnosed that automatic recovery levels are not likely to be successful. In this case, all intermediate levels can be omitted and the final recovery level may be applied. An example is the Lucent Technologies Reliable Clustered Computing (RCC) product [1]. The RCC product incorporates various recovery strategies to guarantee availability of commercial telecommunications systems.

There is a considerable literature on availability estimation in the situation where system parameters are known; our work is motivated by the recent model proposed in [2] and [3]. In this paper we address the situation where the value of some of the system parameters are unknown, but that some knowledge of this uncertainty is available. The uncertainty about parameter values is quantified and subsequently modelled by probability distributions.

The number of parameters of a system can be large and these parameters can be related in complicated ways. In order to predict the availability of a system with variable behaviour, a probability model on the parameters must be defined that takes account of the relationships between parameters. A common technique in stochastic modelling for this situation is the use of graphical models to specify the relationships between variables [4]. A Bayesian Network (BN) is a type of graphical model. The advantages of using BNs is that they facilitate the clear specification of the relationships between system parameters, and they also lead to significant computational savings when implementing statistical inference. The versatility of BNs is becoming increasingly recognised in the field of reliability [5]–[9]. BNs are also seeing increasing use in other fields, such as the modelling of genetics [10] and financial systems [11].

This paper is structured as follows. Section III describes the segregated failure model for availability prediction we discuss and introduces the example system. Section IV models the example system as a random process. Section V introduces BNs, and shows how to simulate

the predictive distribution of the downtime of the system using the BN for the example system. Section VI incorporates the information from observed data into the model, again updating the predictive distribution. Section VII gives examples of the use of the BN for the RCC system as a decision making tool. Section VIII provides a discussion and some concluding remarks.

III. Segregated Failure Model

A. Model description

Denote the set of all possible types of failures of the system by $\boldsymbol{F}$. If a failure occurs in the system, and recovery level j is applied, then we say that the failure was *served* at level j. Each consecutive recovery level can either succeed or fail when applied. Denote the set of failures that are served a level j by $\boldsymbol{F}_j \subseteq \boldsymbol{F}$. The recovery levels are applied consecutively, so $f \in \boldsymbol{F}_j$ implies $\Rightarrow f \in \boldsymbol{F}_l,\ 0 < l < j$.

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. Denote the set of failures of type j by $\boldsymbol{F}_{type_j}$. Then $\boldsymbol{F}_{type_j} \subseteq \boldsymbol{F}_j$. The sets $\boldsymbol{F}_{type_j}$ partition $\boldsymbol{F}$, i.e.,

$$\boldsymbol{F} = \cup_{j=1}^{k} \boldsymbol{F}_{type_j}.$$

and

$$\forall\ j, l : 1 \leq j, l \leq k,$$

$$\boldsymbol{F}_{type_j} \cap \boldsymbol{F}_{type_l} = \emptyset.$$

There are three possible outcomes if a failure is served at level j. The first 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, is denoted p_{rec_j}, i.e.,

$$p_{rec_j} = \text{Prob}[f \in \boldsymbol{F}_{type_j} \mid f \in \boldsymbol{F}_j].$$

The second outcome is that the system is not recovered and no specific hardware failure is identified by the recovery level. In this situation the recovery level $j + 1$ is applied. The conditional probability of applying the next recovery level, given that recovery level j has been applied, is denoted p_{next_j}, i.e.,

$$p_{next_j} = \text{Prob}[f \in \boldsymbol{F}_{j+1} \mid f \in \boldsymbol{F}_j].$$

The third outcome is that the system is not recovered and some specific hardware failure is identified, making manual repair or replacement necessary. Manual repair or replacement are performed at the last recovery level. The probability that the manual repair or replacement recovery level is applied, given that recovery level j has been applied, is denoted p_{last_j}, i.e.,

$$p_{last_j} = 1 - p_{rec_j} - p_{last_j}.$$

Probability p_{rec_j} is defined for $1 \leq j \leq k$, while p_{next_j} and p_{last_j} are only defined 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 $p_{last_k} = p_{next_k} = 0$ and $p_{last_{k-1}} = 0$.

There exist failures that can be escalated to the final recovery level before applying recovery level 1. Denote the probability that a failure will be escalated to the last recovery level before applying level 1 as p_{last_0}, and its complement as $p_{next_0} = 1 - p_{last_0}$.

The rate of all failures is denoted λ. The amount of time it takes to apply recovery level j is denoted τ_j. The reciprocal of τ_j, $\mu_j = 1/\tau_j$, is referred to as the restoration rate.

Fig. 1 shows the diagram of the system and its parameters given in [2].

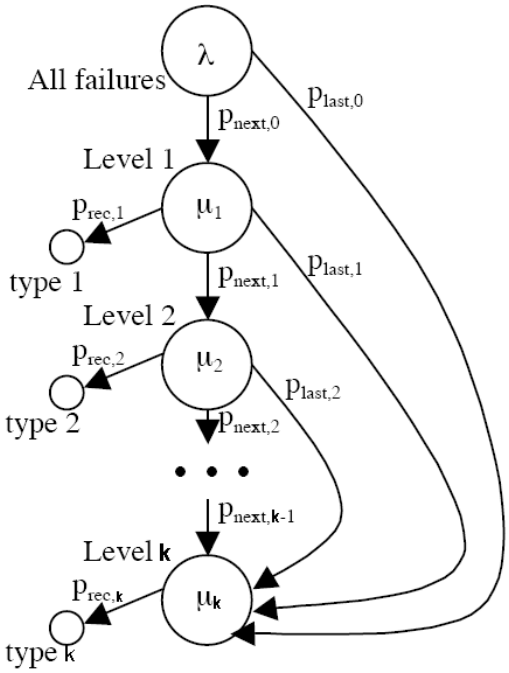

Fig. 1. Classifying failures into k types [2]

Every recovery level is described by $(\mu_j, p_{rec_j}, p_{next_j}, p_{last_j})$. A system with k recovery levels is described by $4k+2$ parameters, $(\lambda, p_{next_0}, (\mu_j, p_{rec_j}, p_{next_j}, p_{last_j})_{j=1}^{k})$.

B. Model Evaluation

The probability that a failure is of type 1 is evaluated by

$$p_{type_1} = p_{rec_1} p_{next_0}. \tag{1}$$

The probability that a failure is of type $j,\ 1 < j < k$, is evaluated as

$$p_{type_j} = p_{rec_j} \left(\prod_{l=1}^{j-1} p_{next_j} \right) p_{next_0}. \tag{2}$$

For the failures that will only be recovered by the final recovery level k, the probability is calculated as

$$p_{type_k} = \left(p_{last_1} + \sum_{j=2}^{k-1} \left\{ p_{last_j} \prod_{l=1}^{j-1} p_{next_l} \right\} + \prod_{l=1}^{k-1} p_{next_l} \right) \tag{3}$$

$$\times\, p_{next_0} + p_{last_0}, \tag{4}$$

or simply

$$p_{type_k} = 1 - \sum_{j=1}^{k-1} p_{type_j}. \tag{5}$$

The rate of all failures is denoted λ_{type_j} and is calculated the product of the overall failure rate λ and the probability of a failure of that type, p_{type_j},i.e.,

$$\lambda_{type_j} = \lambda p_{type_j},\ j = 1, \ldots, k.$$

The amount of time it takes to recover a failure of type j is denoted τ_{type_j}. The service time, denoted $\tau_j,\ j = 1, \ldots, k$, is the time it takes to perform recovery level j. The total recovery time should include the service time, τ_j, and the recovery time spent on the previous recovery levels, i.e.,

$$\tau_{type_j} = \sum_{l=1}^{j} \tau_j.$$

Denote the expected downtime due to failures of type j in a year as t_{d_j}. Thus

$$t_{d_j} = \lambda_{type_j} \tau_{type_j}.$$

Hence, the expected total downtime of the system is $T_d = \sum_{j=1}^{k} t_{d_j}$.

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	0.99
c_1-c_3	0.9

C. Example: Reliable Clustered Computing (RCC)

An RCC system is a collection of connected processors. The sequence of recovery levels considered are:

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.

The parameter values for the model are shown in Table I. The model used in [2] specifies the parameters $p_{last_j} = 0,\ 1 \leq j \leq k$. For this system, serious hardware failures are only diagnosed before applying the procedure of level 1 (i.e., at level 0). 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 the proportion of failures at level j that will be fixed by recovery level j. 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 the proportion of failures that impact all processors in the RCC, requiring a manual repair. Hence,

$$p_{type_1} = c_1 c_a, \quad \text{from (1)}$$

$$p_{type_2} = c_2(1-c_1)c_a \quad \text{and}$$

$$p_{type_3} = c_3(1-c_1)(1-c_2)c_a, \quad \text{from (2)}$$

$$p_{type_4} = ((1-c_1)(1-c_2)(1-c_3)c_a + 1 - c_a), \quad \text{from (3)} .$$

IV. Modelling System Variation

A. Stochastic Model

The downtime of most systems is inherently variable. It may be true that for a given computer system, the same operations will produce the same results every time. Nevertheless, failures are still governed by a random process – the usage profile of the system. Modelling the variability of the system allows us to assign a probability that the system will fail the availability requirements.

We assume that failures occur as a Poisson process with rate λ per year. Failures can only occur when the system is not recovering. We assume that the proportion of time in recovery compared to the total time we are interested in is small, hence the number of failures in a year v is Poisson distributed with mean λ:

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

where $v|\lambda$ is defined to be the distribution of v conditional on knowing the value of λ.

Recovery levels $1, \ldots, k-1$ (the first 3 in our example) are automated, and thus likely to be reasonably stable processes with negligible variance in completion time, and the $\tau_j,\ j = 1, \ldots, k-1$, will be assumed to be constant. The manual recovery level, however, is likely to take a variable length of time to complete. Let t_{k_i} be the i^{th} observed completion time for recovery level k. It will be assumed that given the parameters τ_k_m and τ_k_σ, t_{k_i} follows a Gaussian normal distribution, i.e.,

$$t_{k_i} \mid \tau_k\_m, \tau_k\_\sigma \sim \text{N}(\tau_k\_m, \tau_k\_\sigma^2), \tag{6}$$

truncated at zero. Although the Gaussian distribution is used here, it is trivial to replace this with an alternative distribution.

V. BAYESIAN NETWORKS

A. Graphical Representation of Probability Distribution

Bayesian networks (BNs) are a representation, via a directed acyclic graph, of the dependencies between a set of random variables [12]. The network represents causal relationships between the variables. There is one node for each variable and a directed edge from node A to node B represents that the variable at node A "causes" what happens to the variable at node B; more concretely, the probability distribution of the variable at B is a function of the value of that at node A. A key property of the BN is that the joint probability distribution of all the variables in a network can be decomposed. For a set of n random variables $X_1, \ldots, X_n$ in a network, the joint probability distribution is:

$$p(x_1, \ldots, x_n) \;=\; \prod_{i=1}^{n} p(x_i \,|\, \text{parents of } x_i), \tag{7}$$

where a parent of X_i is a variable whose node has an edge that goes to the X_i node. Fig. 2 shows the Bayesian network of the joint distribution of the subset of the system corresponding to the first recovery level. Rectangular nodes are used to represent constant parameters. This graph assumes that the values of the parameters are known exactly. In practice this will not always be true, and, following the usual Bayesian approach, prior distributions for the parameters about which uncertainty exists must be specified.

VI. INCORPORATING DATA

A. Including Data in the BN

As data are observed, the uncertainty about the parameter values would be expected to decrease. Bayes' theorem provides us with a convenient formula to update our uncertainty upon observation of data. Let v be the predicted number of failures in the system for the next year. Let $(v_1, v_2, \ldots, v_{n_v})$ be observations of the number of failure in a year in n_v systems. Using the structure of the BN, and Equation (7), the joint distribution of $(v_1, v_2, \ldots, v_{n_v})$, v and λ is

$$\begin{aligned} &p(v_1, v_2, \ldots, v_{n_v}, v, \lambda) \\ &\qquad = p(v \mid \lambda) \prod_{i=1}^{n_v} p(v_i \mid \lambda) p(\lambda \mid \lambda\_m, \lambda\_\sigma). \end{aligned}$$

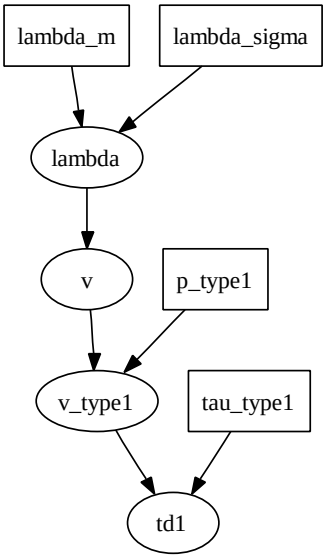

Fig. 2. The BN describing the variability of the subset of the system corresponding to failures of type 1. The BN includes uncertainty about λ, but assumes perfect knowledge about p_{type_1} and τ_{type_1}. The prior parameters are constant and so are represented by rectangular nodes.

Using Bayes' rule, the probability distribution of v given $v_1, \ldots, v_{n_v}$ is then proportional to $\int p(v \mid \lambda) \prod_{i=1}^{n_v} p(v_i \mid \lambda) p(\lambda \mid \lambda\_m, \lambda\_\sigma) d\lambda$.

The BN can be extended to have one node for each observation $i = 1, \ldots, n_v$. A more tidy way to represent the data is to use a shaded rectangular node, symbolising that more than one observation of the same system attribute is represented.

Fig. 3 shows the subgraph with multiple observations. To simulated from this graph, first samples must be drawn from the posterior distribution of λ, then samples can be drawn from $p(v \mid \lambda)$. Depending on the form of the chosen prior distribution and the form of the likelihood function, sometimes the posterior distribution can be found in closed form. If that is not the case, there are alternative methods. For this work, adaptive rejection metropolis sampling (ARMS) was used to simulate from all posterior distributions. This was implemented using the `arms` [13] software package for R [14].

B. Reliable Clustered Computing using the Complete BN

Assume also that n_{τ_k} observations of repair times for the k^{th} level are observed. This example uses four recovery levels, so $k = 4$. The full BN for the RCC system with four recovery levels

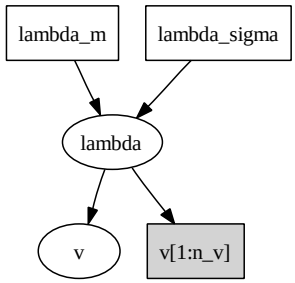

Fig. 3. BN including n_v observations of annual number of failures, $(v_1, v_2, \ldots, v_{n_v})$, and prior uncertainty for λ. The data is represented by a rectangular shaded node to make the graph more tidy.

is given in Fig. 4.

The graph in Fig. 4 can be used to simulate the system, by which we mean simulate from the joint distribution of all the random variables in the system:

$$
\begin{aligned}
& p(c_a, c_1, c_2, c_3, p_{type_1}, p_{type_2}, p_{type_3}, p_{type_4}, \lambda, \lambda\_m, \lambda\_\sigma, \\
& \quad v, v_1, v_2, \ldots, v_{n_v}, v_{type_1}, v_{type_2}, v_{type_3}, v_{type_4}, \\
& \quad \tau_4\_\mu, \tau_4\_\sigma, \tau_{4_1}, \tau_{4_2}, \ldots, \tau_{4_{v_{type_4}}}, t_{d_1}, t_{d_2}, t_{d_3}, t_{d_4}, T_d) \\
& \propto p(p_{type_1} \mid c_a, c_1) p(p_{type_2} \mid c_a, c_1, c_2) \times \\
& \quad p(p_{type_3} \mid c_a, c_1, c_2, c_3) p(p_{type_4} \mid c_a, c_1, c_2, c_3) \times \\
& \quad p(\lambda \mid \lambda\_m, \lambda\_\sigma) p(v \mid \lambda) \times \\
& \quad p(v_{type_1} \mid p_{type_1}, v) p(v_{type_2} \mid p_{type_2}, v) \times \\
& \quad p(v_{type_3} \mid p_{type_3}, v) p(v_{type_4} \mid p_{type_4}, v) \times \\
& \quad p(\tau_4\_\mu \mid \tau_4\_\mu\_m, \tau_4\_\mu\_\sigma) p(\tau_4\_\sigma \mid \tau_4\_\sigma\_m, \tau_4\_\sigma\_\sigma) \times \\
& \quad \left(\prod_{i=1}^{v_{type_4}} p(\tau_{type4_i} \mid \tau_4\_\mu, \tau_4\_\sigma) \right).
\end{aligned}
$$

The procedure for sampling from this joint distribution is to draw samples of the variables that only depend on constants first, and then draw samples from the variables that depend on those, and so on. The only difficulty in this case arises from the fact that τ_4_μ and τ_4_σ are both linked to the constant values $\tau_{4_1}, \tau_{4_2}, \ldots, \tau_{4_{n_{\tau_4}}}$, which makes the posterior values of τ_4_μ

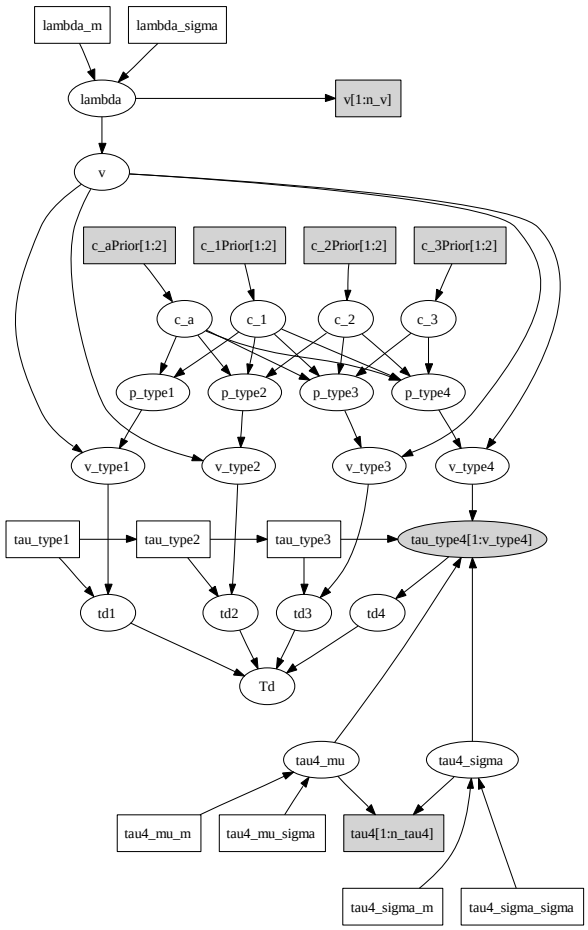

Fig. 4. The full BN for the system. The simulation method to determine the predictive distribution is relatively straightforward to identify once this graph is defined. Notice that the graph does not specify the distributions of individual variables, making it easier to change distributions for different systems or due to new information.

and τ_4_σ dependent. This means that the posterior distribution $p(\tau_4\_\mu, \tau_4\_\sigma)$ is inseparable and both variables must be sampled simultaneously. The approach we used was Gibbs sampling [15], sampling from the conditional distributions $p(\tau_4\_\mu \mid \tau_4\_\sigma)$ and $p(\tau_4\_\mu \mid \tau_4\_\sigma)$ in turn.

Priors for λ, τ_4_μ and τ_4_σ are assumed to be Normal distributions truncated to the positive real line. Prior for the coverage factors c_a, c_1, c_2, and c_3 are assumed to be beta distributions with parameters α_j and β_j, $j = 1, \ldots, k$. The variables $v_{type_1}, \ldots, v_{type_4}$ are assumed to be multinomially distributed given the parameters $p_{type_1}, \ldots, p_{type_4}$ and sample size v. The rest of the variables can be determined functionally from these. The required parameter values for the prior distributions are given in Table II. The values of the prior parameters are 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
$\tau_4\_\mu\_m$	240	c_{1_α}	9
$\tau_4\_\mu\_\sigma$	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	c_{3_α}	99
τ_2	5	c_{3_β}	1
τ_3	7		

To generate data, ten observations of failures per year were simulated from a Poisson distribution with $\lambda = 8$. Also, observations of repair times of recovery level 4 were simulated from a Normal distribution, truncated on the positive real line, with a mean of 240 and a standard deviation of eight.

The predictive distribution of downtime T_d over one year, simulated from $L = 10\,000$ runs, is displayed in Fig. 5. The density functions in Fig. 5 and Fig. 7 are estimated using Gaussian kernels.

The multi-modal shape of the probability density function (pdf) is caused by the length of time it requires to do a manual repair, and each peak corresponds to a different number of manual repairs. The graph is highly skewed; the mean value of T_d is 198 minutes and the median downtime is 20 minutes. This corresponds to a predicted mean availability of 0.9996252 and a predicted median availability of 0.999962.

Further features of the system may be established using the simulation of the joint posterior distribution. For example, it might be useful to know the distribution of failures when the downtime is greater than 300 minutes, in order to determine which failure types are causing the most extreme losses of availability. This can be found by selecting all the simulated samples with $T_d > 300$ and analysing the distribution of the other variables in those samples. Summary statistics are given for the downtimes due to each type of failure in Table III. In some years

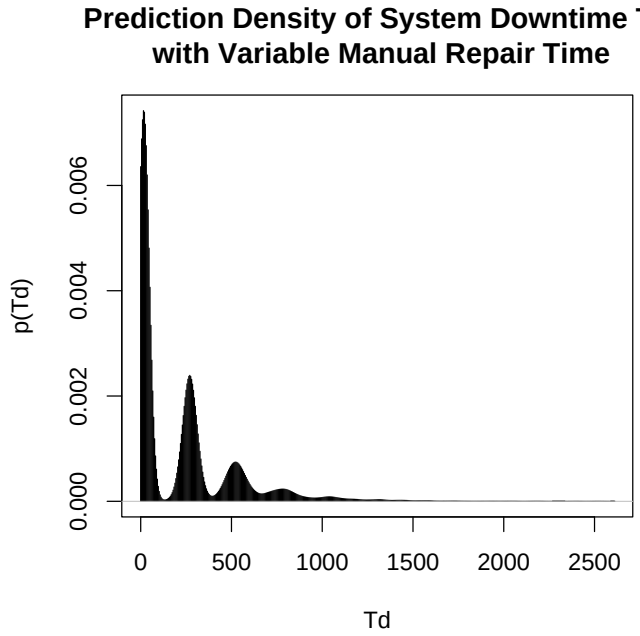

Fig. 5. The predictive distribution for downtime of the system over one year.

no failures occur so the minimum downtime is zero. No failures of types 2, 3, or 4 occur in more than half of the system years so the median values for t_{d_2}, t_{d_3} and t_{d_4} are zero. Summary statistics are given for the downtimes due to each type of failure when downtime T_d is greater than 300 minutes in Table IV. From Tables III and IV it is clear that the main cause of loss of availability is due to the manual repairs.

TABLE III

DISTRIBUTION OF MINUTES/YEAR FOR EACH TYPE OF DOWNTIME.

	2.5% Percentile	Median	Mean	97.5% Percentile
t_{d_1}	2	12	12.85	24
t_{d_2}	0	0	4.473	21
t_{d_3}	0	0	0.977	14
t_{d_4}	0	0	197.3	934

It is important to know not just an estimate of the expected downtime of one system over one year, but an idea of the variation in the predicted distribution. We can calculate prediction

TABLE IV

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

	2.5% Percentile	Median	Mean	97.5% Percentile
t_{d_1}	2	12	12.24	24
t_{d_2}	0	0	4.534	21
t_{d_3}	0	0	1.027	14
t_{d_4}	244.2	528	615	1325

intervals for T_d by ordering the simulated values of T_d and taking percentiles. For example, the fifth and ninety-fifth percentile of the simulated sample give us a 90% prediction interval. For this data, the 90% prediction interval is $[8, 804]$.

VII. DECISION MAKING UNDER UNCERTAINTY

Having the joint probability distribution of all the variables that describe the availability of our system is not necessarily the goal of an analysis. Often, the objective of this type of analysis is to gain the ability to make decisions about the system. Here we use decision theory, the natural partner to Bayesian statistical methods for decision making under uncertainty [16]. This requires us to specify a function, the utility function, that describes the worth of different outcomes of a decision to us. The outcome depends on unknown quantities. We take the action that maximises expected utility. The outcome is often monetary. Here we give two examples.

For a set of outcomes $\boldsymbol{x}$, a utility function is a map $U : \boldsymbol{x} \rightarrow \mathbb{R}$, mapping each outcome to a value on the real line. As an example, consider offering a customer a warranty policy stating penalties to be paid depending on the performance of the system over the first year of use. Suppose if the downtime is greater than 600 minutes a penalty is paid to the customer of $\$200 + (T_d - 600) \times \$10$, if the downtime is between 400 and 600 minutes a penalty of $\$200$ is paid, and if the downtime is less than 400 minutes no penalty is paid. The cost of the warranty policy can be represented by the utility function

$$\begin{aligned} U(T_d) = &- 200(H(T_d - 400) - H(T_d - 600)) \\ &- H(T_d - 600)(200 + 10(T_d - 600)), \end{aligned}$$

where $H(\cdot)$ is the Heaviside function. Fig. 6 shows a picture of the return from the warranty policy plotted against different values of T_d.

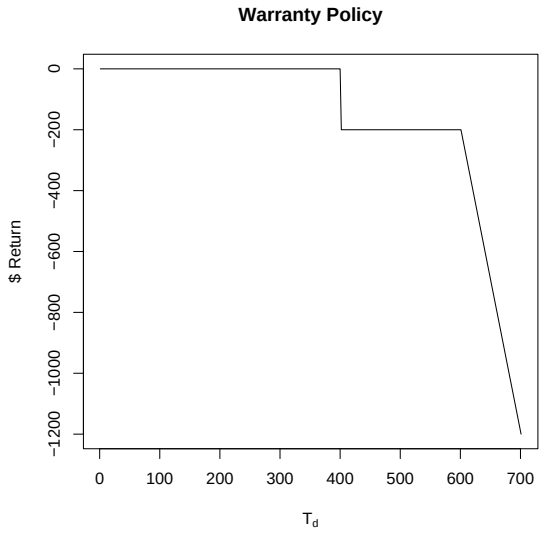

Fig. 6. The warranty policy depends on the observed downtime over the first year of use of the system with the function $U(T_d) = -200(H(T_d - 400) - H(T_d - 600)) - H(T_d - 600)(200 + 10(T_d - 600))$.

The results of the simulation of T_d are used to determine the expected financial cost of the warranty policy after one year. If the simulation contains L samples, the expected cost can be calculated as

$$\frac{\sum_{i=1}^{L} U(T_{d_i})}{L}. \tag{8}$$

Suppose that there is a more expensive way to manually repair the system that always takes 150 minutes and costs $\$100$ per repair. The BN can be adjusted with this information, and, in conjunction with the utility function, can be used to determine if it is of value to always implement this alternative more expensive manual repair. This action is a simplification of the BN, reducing the set of parents of τ_{type_4} to $\{\tau_{type_3}, v_{type_4}\}$.

The simulated pdf of the predicted downtime for the system with fixed manual repair times is displayed in Fig. 7. The multi-modal structure of the predictive distribution is similar to that of the system with variable manual repair time, but the modes are closer together.

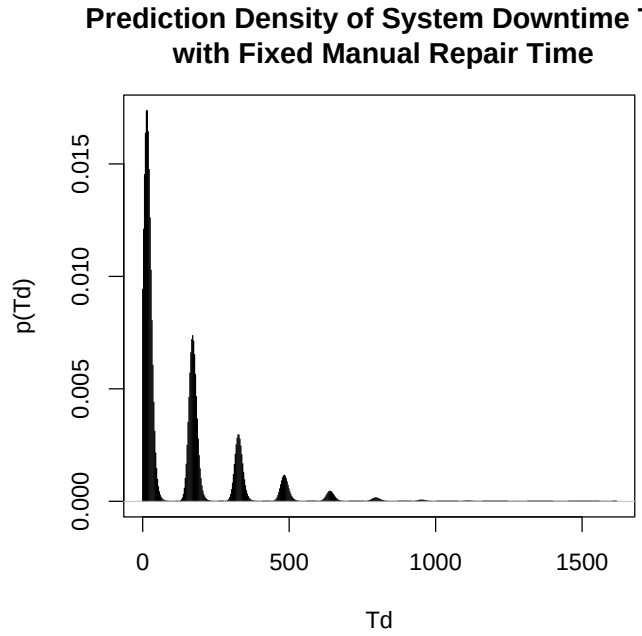

Fig. 7. Histogram of the predicted downtime distribution of the system with fixed manual repair time.

The expected utility of the system with variable repair time is

$$\sum_{i=1}^{L} \frac{U(T_{d_i})}{L} = -\$329.$$

The cost of the system with fixed manual repair times must take into account the costs of the individual repair times. Calculations like this are straightforward from simulated data. The expected financial return for the system with fixed manual repair times is

$$\sum_{i=1}^{L} \frac{U(T_{d_i}) - 100 v_{type_{4_i}}}{L} = -\$40. \tag{9}$$

Therefore, fixing the manual repair time to 150 minutes at an extra cost of $\$100$ per repair gives a reduction in expected cost of the warranty policy of $\$289$ compared to using the variable, less expensive, manual repair.

A. Continuous Decision Space

Another interesting challenge is to optimise the warranty policy. Maintaining the same overall shape of the original warranty function, we can parameterise the utility function using the

parameters p, t_1, t_2 and s as

$$U(T_d, v_{type_4}, p, t_1, t_2, s) = -p(H(T_d - t_1) - H(T_d - t_2)) \\ -H(T_d - t_2)(p + s(T_d - t_2)) - 100v_{type_4}.$$

We wish to maximise the expected utility of a change of warranty over the four-dimensional parameter space of p, t_1, t_2 and s, where $p, t_1, t_2, s \in [0, \inf)$ and $t_1 < t_2$. A new utility function must be set up to take into account the effect on sales that a change of warranty will have. Let $B = 1$ be the event that a customer buys the system and $B = 0$ be the event that a customer does not buy the system. If $B = 1$ then the product is sold for a profit P and incurs a cost $U(T_d, v_{type_4}, p, t_1, t_2, s)$, whereas if $B = 0$ then nothing happens and the profit/loss is 0. The expected utility, taking expectation with respect to B, is then

$$R(T_d, v_{type_4}, P, p, s, t_1, t_2) = \\ p(B = 1)(P + U(T_d, v_{type_4}, p, t_1, t_2, s))$$

Our objective is to maximise the expectation of $R(T_{d_i}, v_{type_{4_i}}, P, p, s, t_1, t_2)$, over the parameters p, t_1, t_2 and s,

$$\begin{aligned} &\mathrm{E}[R(T_{d_i}, v_{type_{4_i}}, P, p, s, t_1, t_2)] \\ &= \mathrm{E}\left[p(B = 1)(P + U(T_{d_i}, v_{type_{4_i}}, p, t_1, t_2, s))\right] \\ &= p(B = 1)\Big(P + \\ &\int U(T_{d_i}, v_{type_{4_i}}, p, t_1, t_2, s))p(T_d, v_{type_4})\, dT_d dv_{type_4}\Big). \end{aligned}$$

The term $\mathrm{E}\left[U(T_{d_i}, v_{type_{4_i}}, p, s, t_1, t_2)\right]$ simplifies to

$$\begin{aligned}
&\mathrm{E}[-p(H(T_d - t_1) - H(T_d - t_2)) \\
&\quad - H(T_d - t_2)(p + s(T_d - t_2)) - 100 v_{type_4}] \\
&= -\mathrm{E}\left[p(H(T_d - t_1) - H(T_d - t_2))\right] \\
&\quad - \mathrm{E}\left[H(T_d - t_2)(p + s(T_d - t_2))\right] \\
&\quad - \mathrm{E}\left[100 v_{type_4}\right] \\
&= p \times \mathrm{p}\left[t_1 < T_d < t_2\right] \\
&\quad - (p - s t_2) \times \mathrm{p}\left[T_d > t_2\right] \\
&\quad - s \times \mathrm{E}\left[T_d H(T_d - t_2)\right] \\
&\quad + 100 \times \mathrm{E}\left[v_{type_4}\right] \\
&= p \times \mathrm{p}\left[t_1 < T_d < t_2\right] \\
&\quad - (p - s t_2) \times \mathrm{p}\left[T_d > t_2\right] \\
&\quad - s \times \mathrm{E}\left[T_d | T_d > t_2\right] \\
&\quad + 100 \times \mathrm{E}\left[v_{type_4}\right].
\end{aligned}$$

We assume here that $p(B = 1)$ will take a logistical shape of the form

$$p(B = 1) = \frac{1}{1 + e^{-\alpha - \beta_1 p - \beta_2 t_1 - \beta_3 t_2 - \beta_4 s}}.$$

Data from customer surveys or past sales experience with warranties will allow estimation of α and the β_i.

Hence, we must find

$$\begin{aligned}
\arg \max_{p, t_1, t_2, s} \Bigg(& \frac{1}{1 + e^{-P + Q(\mathrm{E}[T_d], p, t_1, t_2, s)}} \\
& \times \Big(p \times \mathrm{p}\left[t_1 < T_d < t_2\right] \\
& - (p - s t_2) \times \mathrm{p}\left[T_d > t_2\right] - s \times \mathrm{E}\left[T_d | T_d > t_2\right] \\
& + 100 \times \mathrm{E}\left[v_{type_4}\right] \Big) \Bigg).
\end{aligned}$$

We do not have to re-run the simulation of the BN for the optimisation problem. Let $P = 1000$, $\alpha = 2$, $\beta_1 = 0.000008$, $\beta_2 = \beta_3 = -0.00175$ and $\beta_4 = 0.014$. Searching the parameter space of p, t_1, t_2 and s, using the `optim` function in R, we get

$$
\begin{aligned}
&p = 263.33 \\
&t_1 = 244.77 \\
&t_2 = 244.77 \\
&s = 18.12 \\
&S(p, t_1, t_2, s) = 0.09 \\
&\mathrm{E}\left[U(T_{d_i}, v_{type_{4_i}}, p, s, t_1, t_2)\right] = -265.36
\end{aligned}
$$

Without incorporating uncertainty into our model we would only get an estimate of expected downtime, rather than the predictive distribution of downtime. The expected downtime for the system, calculated as the mean of the predicted downtimes T_d, is equal to 215 minutes, which has zero penalty under the warranty policy. The expected cost of the downtime is greater than the cost of the expected downtime because the former weights downtimes by importance. This is a clear advantage of the modelling approach used here.

VIII. Discussion and Conclusions

It is a common practice in mathematically modelling system to return a constant output parameter for different values of the input parameters. In this paper we used prior distributions to fully specify our knowledge about the state of each parameter before calculating the model. It is then shown how this knowledge can be updated using available data.

The approach to applying the BN used here was based on simulation rather than the explicit calculation of joint probability distributions. In practice, as the complexity of the system being modelled increases, and as it is desirable to model the behaviour of the system as closely as possible, it is rare that explicit formulae for the joint probability distribution can be calculated. For more complex systems, the speed of the simulation can become important and it becomes useful to spend more effort choosing prior distributions that fit together mathematically. For evaluating the comparative utility of a large set of actions, however, it is always possible to run the network for each outcome in parallel on different machines.

The disadvantages of our approach is that it is dependent on information from domain experts whose time is generally valuable, and that the computation time for simulations may become

expensive. It can also be difficult to specify a utility function. However, the coherent aggregation of all available information and expertise is the only way to optimally make business decisions.

The BN approach offers an elegantly straightforward method of propagating uncertainty due to system variation, uncertainty due to lack of information, and information due to expert knowledge and data through a mathematical model. As has been shown here, the bulk of the work in simulating the joint probability distribution of the variables that model a system is in specifying the pairwise links between the variables, and the local conditional distributions. A possible extension to the approach in this paper would be to use continuous-time Bayesian networks (CTBNs) [5]. Further extensions will be made to incorporate cost of data collection and continuous decision spaces, potentially using the method of sampling the decision space in [17].

REFERENCES

[1] G. Hughes-Fenchel, "A flexible clustered approach to high availability," in *27th Annual International Symposium on Fault-Tolerant Computing, FTCS-27*, Seattle, Washington, Jun.

[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*, 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] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988.

[13] G. Petris, L. Tardella, and W. Gilks., *HI: Simulation from distributions supported by nested hyperplanes: R package version 0.1-1*, 2005.

[14] 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

[15] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis, Second Edition*. Chapman & Hall/CRC, 2003.

[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.