Reliability Engineering and System Safety 56 (1997) 29–42

ELSEVIER PII: S0951-8320(96)00134-2
0951-8320/97/$17.00

Using higher-level failure data in fault tree quantification

Harry F. Martz[a] & Russell G. Almond[b]

[a] *Technology & Safety Assessment Division, Los Alamos National Laboratory, Los Alamos, NM 87545 USA*
[b] *Educational Testing Services, Princeton, NJ 08541, USA*

(Received 8 November 1996)

This paper presents a Bayesian method which can simultaneously combine basic event and statistically independent higher event-level failure data in fault tree quantification. Such higher-level data could correspond to train, subsystem or system failure events. In fact, because highest-level data are usually available for existing facilities, the method presented here allows such data to be propagated to lower levels. The method has two stages: (1) a top-down propagation scheme which allocates the higher event-level information to the basic events, at a cost of making them dependent; and (2) a scheme for sampling the probabilities of the dependent basic events. A simple example illustrates the performance of the method. © 1997 Elsevier Science Limited.

1 INTRODUCTION

Vesely *et al.* [1], the probabilistic risk assessment (PRA) procedures guide [2], and many other textbooks, discuss fault tree quantification. Such quantification consists of three steps: (1) determining the basic event probabilities, (2) calculating the minimal cut set probabilities, and (3) determining the system (i.e., the top event) probability. A simple modification of the procedure produces a quantitative measure of importance of the basic events, and some of these importance measures are related to the probability of system failure conditioned on the occurrence of a basic event.

It is current and accepted practice in fault tree and accident sequence quantification [as implemented, for example, in the Systems Analysis Programs for Hands-on Integrated Reliability Evaluations (SAPHIRE) [3] package Integrated Reliability and Risk Analysis System (IRRAS) [4, 5] program] to use only statistical data and information regarding the basic events. Furthermore, it is presently impossible to directly use independent statistical data and information corresponding to higher-level events or gates in the tree, despite the fact that normal operation and testing procedures often generate these data for many high-level gates including those corresponding to such events as train, subsystem, and system unavailability, and even the top event itself. In quantifying the accident sequence frequency for a proposed accident of interest at an existing facility, independent statistical data almost always exist at the highest level; namely, x occurrences of the accident (where x is often 0) in a given exposure time t or in n demands.

By 'independent' we mean that the higher-level data for a system are *not* simultaneously providing collateral information on the basic events comprising that system (which would lead to double counting and thus dependency). In other words, we assume that the higher-level and any basic-event data are *not* the result of the same set of demands or observation period. This is usually the case for any system test that is destructive, such as a missile fired at a target. If the same higher-level data provide basic event-level information, then we can instead use such data to verify the structure of the fault tree. In particular, any higher-level failure data which is not predicted by the fault tree is an indication that the fault tree model is inadequate.

In addition to higher-level statistical test data, independent industry-wide statistical analyses are sometimes performed on safety systems considered in a PRA. Such analyses represent a source of generic higher-level statistical information for the specific plant under consideration. For example, Grant *et al.* [6] describe an industry-wide statistical analysis of the safety-related performance of the high-pressure coolant injection (HPCI) system at US commercial boiling water reactor plants for the period 1987–1993. If we had a fault tree quantification method capable of utilizing such higher-level information, then the results of such analyses effectively represent an additional independent source of statistical information supplementary to any plant-specific basic event data.

This paper describes a Bayesian method which can simultaneously combine basic event and independent higher-level failure data and information in fault tree quantification. The obvious advantage is the associated increase in accuracy and precision of the probabilistic results because of the combined use of these data. However, there is another important but somewhat less obvious advantage. The method will provide us with the opportunity to compare two independent estimates for the gate-level probability: one based on the gate-level data and one induced by the logical model of the gate and the associated basic event and/or gate-level data below it in the tree. Such a comparison represents a direct quantitative evaluation of the adequacy of the logical model and data for the gate in question; thus, the method is useful for model validation.

1.1 Related methods

Several authors consider the combined use of both component and independent system-level test data in reliability analysis. Easterling and Prairie [7] develop classical maximum likelihood estimators of series or parallel system reliability using both types of data; however, their method is quite restrictive in that the components must be both independent and identical. Mastran [8] considers Bayesian estimation of system reliability in which there exist test data at both the component and system levels for both binomial pass/fail and exponential time-to-failure data for a series system of nonidentical components. He uses a top-down approach which apportions the posterior system reliability distribution to each component in the form of a component prior distribution consistent with the series configuration. Combining these component priors with the component level data (using Bayes' theorem) produces component posterior distributions. Propagating these component posteriors back up to the system level using the series model forms the final system posterior from which the desired inferences are obtained. Mastran and Singpurwalla [9] extend the method to include any coherent system, and Barlow [10] likewise also considers a Bayesian method for combining both types of data.

Martz *et al.* [11] develop a Bayesian procedure for estimating series system reliability that permits the use of both types of data at the component, train, subsystem, and system levels. Martz and Waller [12] extend the method to include an arbitrarily complex system configuration of series/parallel subsystems of other subsystems or components. Martz and Baggerly [13] further extended the method to accommodate Poisson as well as binomial data at any level in a complex series/parallel system.

Hulting and Robinson [14] consider pass/fail data, lifetimes of nonrepairable components, and repair histories for repairable subsystems at any level in a series system. They consider a censored Weibull model for the pass/fail data, a Weibull model for the lifetime data, and a power law process model for the repair data. They also express informative component prior information by means of a conjugate weighted-gamma family of prior distributions. They use their method to estimate the reliability of a new automobile vehicle system in the early stages of development.

When both levels of data exist for the same demands or observation period, the above methods are inapplicable because the data are dependent. For example, a standby system may fail to operate upon demand (a higher-level system failure) which may subsequently be traced to the failure of a particular component in the system (a basic event-level failure). However, the above methods (and the method presented here as well) are still applicable if only one level of data is used. Using the data at the higher-level gate to form an aggregated posterior for the higher-level event produces an *aggregate* analysis. Using the data at the basic event-level to form a disaggregated posterior for the higher-level event produces a *disaggregate* analysis. Usually, the aggregate and disaggregate posteriors will disagree, in which case we say that an *aggregation error* occurs. Very large aggregation errors are grounds for suspicion of the structure of the fault tree model.

Bier [15] and Azaiez and Bier [16, 17] develop necessary and sufficient conditions (which are extremely stringent) for perfect aggregation (that is, no aggregation error) for several classes of reliability models. They also demonstrate why analysts must be careful in choosing the level at which to use such dependent data and argue that, when accurate disaggregate data, prior distributions, and reliability models are available, concerns about aggregation error favour using a disaggregate analysis. Bier [15] and Azaiez and Bier [16, 17] provide additional information and insight regarding this important decision.

1.2 Brief overview of the method

All of the above methods place restrictions on either the structure of the system or the distribution of the probabilities of the basic events. In Mastran and Singpurwalla, [9] this restriction is the subtle one that component level information can be expressed by independent data. In contrast, the method presented here does not rely on specific distributional assumptions regarding the basic or higher event-level probabilities. We assume that the information about the probability of occurrence of each basic event can be summarized with a probability distribution with known mean and variance from which we can draw values. The method describe here will even properly handle *a priori* state-of-knowledge (SOK) dependence among the basic events; all that is necessary is a method for sampling their joint distribution. Similarly, the method requires that the higher event-level data or information reside in the form of a distribution of the probability of occurrence of the higher event with known mean and variance from which we can likewise randomly sample. In the case of higher-level data, this could be a posterior distribution derived from either an informative or noninformative prior.

The primary part of the method described in this paper is an algorithm for allocating the higher-level information to the basic events according to the importance of each basic event. The allocation algorithm provides both a mean and a variance for the allocated information. The final basic event probability is a mixture of the original and allocated event information, using the precisions of the basic event probabilities (inverse variances) as the mixing weights. Propagating the basic event mean occurrence probabilities up to the higher event-level gate produces a new gate-level mean occurrence probability.

The higher-level gate mean occurrence probability is a biased estimate because it does not take into account SOK dependence among the basic events. Even if the original basic event distributions were *a priori* SOK independent (for example, no common correlation classes), the allocation algorithm induces a SOK dependence among the basic events. This is easy to see because, after the allocation, the knowledge about each basic event depends on the common information at the higher-level gate.

The second part of this method is an algorithm for sampling from the joint distribution of basic event probabilities which accounts for the SOK dependence. It consists of sampling the basic event probabilities in the normal way and adding a simple adjustment for the higher event-level information. This second algorithm is used in a simple Monte Carlo procedure to produce unbiased estimates of the higher-level gate mean occurrence probability.

A model for using independent higher-level failure data in any coherent fault tree is presented in Section 2. Section 3 illustrates the performance of the method using a simple fault tree example, and Section 4 discusses the strengths and limitations of the procedure based on our experience with it.

2 A MODEL FOR USING HIGHER-LEVEL FAILURE DATA

We consider a top-down approach which apportions the higher-level data down the fault tree to update the basic event probabilities so that the *combined* basic event probabilities reflect both their initial values and the higher-level data. In this sense, the higher-level data *induce* changes in the basic event probabilities that reflect the higher-level data. This approach implicitly assumes that the structure of the fault tree is correct. Even in cases where this assumption does not hold, the combined basic event probabilities should produce better estimates of the higher-level gate occurrence probability (in the sense that it will be more consistent with the higher-level data) although the combined basic event probabilities will not necessarily be better estimates of the individual basic event probabilities.

While a corresponding bottom-up approach, in which the probabilities of the higher-level events would be updated, could be used, there are several reasons why we prefer the top-down approach. First, a top-down approach requires that event data be maintained (and subsequently updated) only at the basic event-level of the tree, a procedure that is consistent with all existing fault tree quantification software. Thus, the top-down approach is relatively easy to integrate into existing software. Second, because the basic event probabilities ultimately reflect all of the combined data, these probabilities dynamically represent the most current SOK available. Because such data are uniformly maintained at the lowest possible level, it can be directly used to estimate the unavailability of other systems, including those not appearing in the current fault tree. Third, in those cases in which informative subjective assessments of system unavailability are to be incorporated at only one level in the analysis (that is, informative prior distributions are to be considered at either the higher or basic event-level of the tree—but not both), we believe that such priors are more meaningfully assessed at the higher event-level. Mastran [8] and Mastran and Singpurwalla [9] agree and likewise consider a top-down approach.

Once we have updated the basic event probabilities to reflect the higher-level data at a given gate, we can then propagate these combined probabilities back up the fault tree to update the gate probability of occurrence. This probability estimate likewise reflects both the initial basic event-level data as well as the higher-level data at the gate. We then iterate this procedure, sequentially updating the basic event-level probabilities, until all the gates for which there exist higher-level data have been considered. Finally, we then propagate the final updated basic event-level probabilities through the fault tree to obtain the required estimate of the higher-level gate probability of occurrence. The details for accomplishing this are described below.

This procedure requires that there exists a posterior SOK *uncertainty distribution* of the probability of each basic event involved in the updating process (that is, those basic events contributing to the probability of the higher-level gate). This is a Bayesian approach based on the assumption that SOK uncertainty distributions of the parameter(s) of the corresponding model induce a corresponding SOK uncertainty distribution on the probability of occurrence of the basic event. This posterior distribution may be computed directly from the posterior distributions on the model parameters which are required when performing an uncertainty analysis. We further assume that the SOK posterior distribution for each basic event has a known mean and variance.

The procedure similarly requires that the failure information at the higher-level gate be expressed as an *independent* SOK probability distribution with known mean and variance on the higher-level gate probability of occurrence. For example, if the higher-level information consisted only of plant-specific statistical surveillance and/or operating performance data, the higher-level information would be used to construct a posterior distribution using a noninformative prior. If the independent information is the result of an independent system analysis (such as an industry-wide system performance analysis as in Grant *et al.* [4]), then it would be summarized as an informative prior distribution for the probability of the gate. If both types of information are available, the independent analysis could be used to determine an informative prior and the resulting posterior would become the SOK probability distribution for the gate. All that is required is that the information sources used in forming the SOK distribution for the higher-level gate occurrence probability be *independent* of the information used to form the SOK distributions of the probabilities of the basic events. Any type of probability model, such as those based on the binomial or Poisson distribution, can be used, provided that the model used ultimately produces a probability as output. As in the case of the basic events, we require the final SOK posterior distribution of the gate probability along with its mean and variance.

The use of higher-level data usually induces a SOK dependence between the basic events affected by the higher-level data. This dependence is a consequence of using the common higher-level data by all the basic events affected by the higher-level data. However, cases exist in which the common use of dependent data doesn't necessarily induce SOK dependency (see Haim [18]). Thus, depending upon the data, the SOK dependency structure, and the importance of the basic events to the gate for which the higher-level data exist, such dependency may or may not be important to consider. However, to the extent that these dependencies exist, the method to be presented recognizes and preserves such induced stochastic dependence among the basic events. Apostolakis and Moieni [19] also discuss SOK dependence in PRA, particularly with regard to common cause failures.

If the same data set is used for a group of similar basic events (or components), then this common usage induces a SOK dependence (or correlation) between the basic events in the group. For example, suppose that a plant has two motor-driven residual heat removal pumps. These pumps are virtually identical, and therefore are modeled as having the same probability of failure to start on demand p in the model. Often, the same data sources contribute to the uncertainty distribution for p for each pump, yielding an identical SOK distribution for p for each pump. In this case, a SOK dependency exists between the two pump failure basic events, and the SOK uncertainty distributions for the two values of p are perfectly correlated. This SOK dependency must be distinguished from basic events in which such SOK dependency does not exist. It is common PRA practice to accommodate such SOK dependency by using the same uncertainty distribution for the group, a practice known as defining a *correlation class*. When a probability is sampled from its SOK uncertainty distribution, that same value is assigned to all the basic events in the class. The method presented below recognizes and permits the use of correlation classes.

Before presenting the method, we introduce some general notation and assumptions. We denote the basic events in the fault tree as Bj; thus, j is used as a basic event index. We denote the gates at which we have higher-level data or information as Si, $i = 1, 2, ..., m$, where $m \geq 1$; thus, i is a gate index. Two subscript styles denote gates and basic events: 'Si' refers to gate Si, and 'j' (without a capital letter) refers to basic event Bj. Each basic event and gate has a corresponding random event which we also denote as Bj and Si. Thus, $P(Bj) = p_j$ and $P(Si) = p_{Si}$ indicate the probability of the basic event Bj occurring and a pattern of basic events occurring such that the gate Si occurs, respectively.

The method given below absorbs the higher-level data one gate at a time, and the ordering of the gates Si, $i = 1, 2, ..., m$, is the order in which the higher-level data are absorbed. We note here that the procedure works no matter what ordering is chosen; however, we recommend that the data at the lower-level gates be absorbed first.

When absorbing the higher-level information at gate Si, the method uses two different types of random variables—those random variables produced via a mixture combination at gate Si (representing *cumulative* information) and those random variables not involving a mixture combination at gate Si (representing *incremental* information). We use the superscript '$*(i)$' to denote random variables based on mixture-combination involving the higher-level data at gate Si and the superscript '$+(i)$' to denote random variables that are *not* mixture-combinations involving the higher-level data at gate Si. For example, $p_j^{*(i)}$ denotes the combined (cumulative) probability of basic event Bj at gate Si which reflects the initial basic event data as well as all the higher-level data from gates $S1$, $S2$, ..., Si, while $p_{Si}^{+(i)}$ denotes the non mixture-combined (incremental) probability of gate Si based exclusively on the higher-level data at gate Si. Distinguishing between these two types will become more clear when the details of the model are presented below.

In addition to the fault tree structure, the method assumes the existence of the following random variables: (1) the initial (before absorbing higher-level data) probability $p_j^{*(0)}$ that Bj occurs (these may be dependent random variables if there is an initial SOK dependence; for example, that due to common correlation classes), and (2) at least one independent higher gate-level probability $p_{Si}^{+(i)}$. The assumption of independent higher-level probabilities says that $p_{Si}^{+(i)}$'s are mutually independent, and that these are all mutually independent of the $p_j^{*(0)}$'s.

The approximate probability of gate Si may be calculated from the vector of individual basic event probabilities, **p**, using any number of exact or approximation methods such as the *rare event* approximation [3], the *minimal cut set upper bound* approximation [3], or a direct exact calculation scheme. Note that the first two methods require the minimal cut sets, which has some computational advantage because we can eliminate any event which does not appear in a minimal cut set from consideration. Let $Si[\mathbf{p}]$ denote the approximate probability of gate Si calculated using the most appropriate of these two approximations. The procedure below evaluates $Si[\bullet]$ at randomly sampled as well as mean values of **p**. The method described below also requires the probability of gate Si conditional on the occurrence of basic event Bj, or equivalently, conditional on $p_j = 1$. We denote this conditional probability as $Si[\mathbf{p}|1_j]$.

Many importance measures are available for assessing the influence (or importance) of each basic event to the approximate gate probability. One of these is the *risk increase ratio* (RIR [3]). RIR is an indication of how much the approximate gate probability would increase if the basic event occurred with probability 1.0 (that is, if the corresponding component failed). RIR is determined by calculating the approximate gate probability with the basic event probability set equal to 1.0 and dividing this quantity by the approximate gate probability calculated with the basic event probability set to its true value. In equation form, the RIR for basic event Bj at gate Si based on the vector $\mathbf{p}^{*(i)}$, with elements $p_j^{*(i)}$, is denoted by $\mathrm{RIR}_j^{*(i)}$ and is given by

$$\mathrm{RIR}_j^{*(i)} = P(Si|Bj)/P(Si) = Si[\mathbf{p}^{*(i)}|1_j]/Si[\mathbf{p}^{*(i)}], \quad (1)$$

where $\mathbf{p}^{*(i)}|1_j$ denotes the vector $\mathbf{p}^{*(i)}$ in which the probability of basic event Bj is set to 1.0. We will have occasion to evaluate eqn (1) both for a random sampled vector $\mathbf{p}^{*(i)}$ and its mean $E[\mathbf{p}^{*(i)}]$. The choice of RIR in the analysis below is not based on its use as a quantitative measure of importance of Bj at gate Si, rather, it provides a convenient method (already supported by many fault tree analysis computer programs) of calculating $P(Si|Bj)$.

One final piece of notation is necessary to describe the Monte Carlo sampling used in the algorithms. The population expected value (mean) and variance of a random variable are denoted by $E(\bullet)$ and $\mathrm{Var}(\bullet)$, respectively, while the covariance between two random variables is denoted by $\mathrm{Cov}(\bullet,\bullet)$. The (Monte Carlo) sample mean and variance of a set of data are denoted by $\mathrm{avg}(\bullet)$ and $\mathrm{samvar}(\bullet)$, respectively, while the sample covariance between two sets of data is denoted by $\mathrm{samcov}(\bullet,\bullet)$. The subscript '$n$' on a random variable, in addition to a subscript 'j' or 'Si', denotes a Monte Carlo sample value of the corresponding random variable for the nth simulation cycle.

2.1 Back-estimation of Bj

The top-down allocation of the higher-level probability of gate Si to the probabilities of the basic events Bj affected by this gate is a key feature of the model, and we refer to this as the *back-estimation* of Bj. Let $\overline{Si}$ denote the event that Si does not occur. Prior to using the higher-level data at gate Si, we define

$$I_j \equiv P(Bj|Si) = P(Si,Bj)/P(Si) = \mathrm{RIR}_j^{*(i-1)}P(Bj).$$

Similarly, we define

$$I_j' \equiv P(Bj|\overline{Si}) = P(\overline{Si},Bj)/P(\overline{Si}) = [1 - \mathrm{RIR}_j^{*(i-1)}P(Si)]P(Bj)/[1 - P(Si)].$$

By the law of total probability, we now consider I_j and I_j' as weights and apply these to the probability of occurrence of gate Si based only on the higher-level data, which we denote by $p_{Si}^{+(i)}$, and its complement as follows:

$$\begin{aligned} p_j^{+(i)} &= P(Bj) = P(Bj|Si)P(Si) + P(Bj|\overline{Si})P(\overline{Si}) \\ &= I_j p_{Si}^{+(i)} + I_j'[1 - p_{Si}^{+(i)}] \\ &= p_j^{*(i-1)}[a_j^{(i)} + b_j^{(i)} p_{Si}^{+(i)}] \end{aligned} \quad (2)$$

where we have defined

$$\begin{aligned} a_j^{(i)} &\equiv [1 - p_{Si}^{*(i-1)}\mathrm{RIR}_j^{*(i-1)}]/[1 - p_{Si}^{*(i-1)}], \\ b_j^{(i)} &\equiv \mathrm{RIR}_j^{*(i-1)} - a_j^{(i)}. \end{aligned} \quad (3)$$

Equation (2) allocates (in a top-down manner) the information in $p_{Si}^{+(i)}$ to the basic event Bj. The probability $p_j^{+(i)}$ on the left-hand side of the equation denotes the information about Bj inherent in p_{Si}^{+i} (given the weights I_j and I_j'). Note that (2) uses $p_{Si}^{*(i-1)} = Si[\mathbf{p}^{*(i-1)}]$ and $p_j^{*(i-1)}$, respectively, to estimate $P(Si)$ and $P(Bj)$. Thus, $p_j^{+(i)}$ and $p_j^{*(i-1)}$ are dependent estimates of the probability of Bj; this dependence is considered below.

Three bounding cases provide insight into the behaviour of this apportionment scheme. First, if basic event Bj does not appear in any minimal cut set for gate Si, then $\mathrm{RIR}_j^{*(i-1)} = 1.0$, and we see from eqn (2) that $p_j^{+(i)} = p_j^{*(i-1)}$; that is, $p_j^{*(i-1)}$ is unaffected by the higher-level data. Second, if Bj is a minimal cut set for gate Si, then $\mathrm{RIR}_j^{*(i-1)} = 1.0/p_{Si}^{*(i-1)}$, and again using eqn (2), we have $p_j^{+(i)} = p_j^{*(i-1)} p_{Si}^{+(i)}/p_{Si}^{*(i-1)}$. In this case, $p_j^{*(i-1)}$ is modified proportional to the ratio of the probability of Si based on the higher-level data to the probability of Si predicted from $\mathbf{p}^{*(i-1)}$. Third, if the probability of Si using the higher-level data agrees with the probability of Si predicted from $\mathbf{p}^{*(i-1)}$; that is, if $p_{Si}^{+(i)} = p_{Si}^{*(i-1)}$, then we again see that $p_j^{+(i)} = p_j^{*(i-1)}$.

To further understand the performance of eqn (2), Fig. 1 gives a parametric plot of the ratio $R \equiv p_j^{+(i)}/p_j^{*(i-1)}$ as a linear function of the ratio $S \equiv p_{Si}^{+(i)}/p_{Si}^{*(i-1)}$ for $\mathrm{RIR}_j^{*(i-1)} = 1$, 20, 40, 60, 80, and 100. Also, we have set $p_{Si}^{*(i-1)} = 0.01$ for all the lines in Fig. 1. In addition to the above bounding cases, we see that there is a proportionately greater change in $p_j^{+(i)}$ relative to $p_j^{*(i-1)}$ as $\mathrm{RIR}_j^{*(i-1)}$ increases for Bj (that is, as Bj becomes more important) for a given value of S. This supports our intuition. Finally, as we decrease $p_{Si}^{*(i-1)}$ by successive orders of magnitude and simultaneously increase both $\mathrm{RIR}_j^{*(i-1)}$ and S by this same factor, R simply changes scale by this factor.

Similarly, Fig. 2 gives a parametric plot of R as linear function of $\mathrm{RIR}_j^{*(i-1)}$ for values of $S = 0.1$, 0.5, 1, 5, and 10, where $p_{Si}^{*(i-1)} = 0.01$ is again used in constructing Fig. 2. We see that, for a fixed importance of Bj, R increases proportionally as S increases. If we decrease $p_{Si}^{*(i-1)}$ by successive orders of magnitude, simultaneously increase $RIR_j^{*(i-1)}$ by the same factor, and keep the same parametric values of S as in Fig. 2, then R remains the same as in Fig. 2. That is, as long as $\mathrm{RIR}_j^{*(i-1)}$ increases and $p_{Si}^{*(i-1)}$ decreases by the same factor, there is no change in R provided the ratio S remains constant.

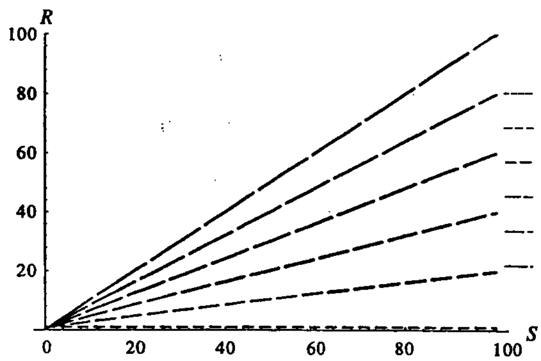

Fig. 1. The ratio R as a linear function of the ratio S for selected values of RIR.

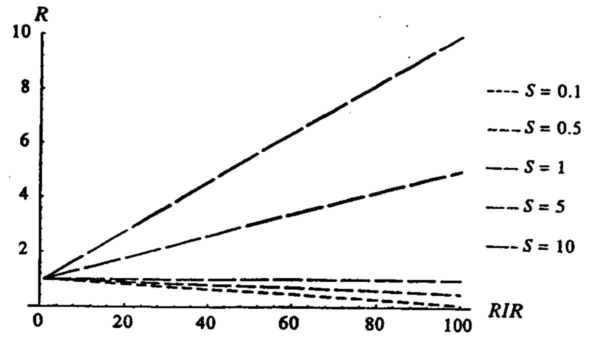

Fig. 2. The ratio R as a linear function of RIR for selected values of the ratio S.

2.2 Combined probability that Bj occurs

The probability $p_j^{+(i)}$ in eqn (2) is the portion of the higher-level event probability $p_{Si}^{+(i)}$ which has been allocated to basic event Bj at gate Si. In order to obtain an updated probability of basic event Bj which reflects all the data at gate Si, we need to combine $p_j^{+(i)}$ with the previously combined probability of Bj from gate $S(i-1)$, $p_j^{*(i-1)}$. A commonly used procedure for accomplishing this is to form a mixture (a weighted average) of both random variables with weights inversely proportional to their respective variances.

Define the *precision*, τ, of a random variable to be the reciprocal of its variance; thus, $\tau_j^{+(i)} = 1/\mathrm{Var}[p_j^{+(i)}]$. The weighted average for combining $p_j^{*(i-1)}$ and $p_j^{+(i)}$ is the random variable given by

$$p_j^{*(i)} = [1 - w_j^{(i)}] p_j^{*(i-1)} + w_j^{(i)} p_j^{+(i)}, \tag{4}$$

where the weight $w_j^{(i)}$ is defined as

$$w_j^{(i)} \equiv \tau_j^{+(i)}/[\tau_j^{*(i-1)} + \tau_j^{+(i)}]. \tag{5}$$

The mean and variance of eqn (4) are

$$E[p_j^{*(i)}] = [1 - w_j^{(i)}] E[p_j^{*(i-1)}] + w_j^{(i)} E[p_j^{+(i)}] \tag{6}$$

and

$$\mathrm{Var}[p_j^{*(i)}] = [1 - w_j^{(i)}]^2 \mathrm{Var}[p_j^{*(i-1)}] + [w_j^{(i)}]^2 \mathrm{Var}[p_j^{+(i)}] + 2w_j^{(i)}[1 - w_j^{(i)}] \mathrm{Cov}[p_j^{*(i-1)}, p_j^{+(i)}] \equiv 1/\tau_j^{*(i)}. \tag{7}$$

Note that $p_j^{*(i-1)}$ and $p_j^{+(i)}$ are dependent by virtue of eqn (2)—hence the covariance term in eqn (7). Also note that, if $p_j^{*(i-1)}$ and $p_{Si}^{+(i)}$ were independent, eqns (6) and (7) would be the Bayesian posterior mean and variance when estimating the mean of a Gaussian distribution with a Gaussian prior where $E[p_j^{+(i)}]$ and $E[p_j^{*(i-1)}]$ are the sample and prior means, respectively.

2.3 Mean, variance and covariance of $p_{Si}^{+(i)}$

Equations (4)–(7) require estimates of $E[p_j^{+(i)}]$, $\mathrm{Var}[p_j^{+(i)}]$ and $\mathrm{Cov}[p_j^{*(i-1)}, p_j^{+(i)}]$. Because eqn (2) is an extremely complex function of all the random variables $p_j^{*(i-1)}$, $j=1,2,\ldots$, it is virtually impossible to analytically determine these moments. Further, because $p_j^{*(i-1)}$, $j=1,2,\ldots$, are correlated, it is extremely complicated to approximate these moments using the usual error propagation formulas for a mean and variance of a general function of several random variables discussed in Section 2.4 below.

Because of these difficulties, we use Monte Carlo simulation to determine $E[p_j^{+(i)}]$, $\mathrm{Var}[p_j^{+(i)}]$ and $\mathrm{Cov}[p_j^{*(i-1)}, p_j^{+(i)}]$. The procedure absorbs the higher-level information one gate at a time, at each time for gate Si for $i=1,\ldots,m$. At each step, the procedure requires only the ability to simulate from the joint SOK distribution of $\mathbf{p}^{*(i-1)}$. For $\mathbf{p}^{*(o)}$ this is just the ordinary simulation which would be used in an uncertainty analysis of the fault tree (before absorbing the higher-level data). For $i>1$, Section 2.5 describes the simulation procedure:

Step 1: Calculate minimal cut sets for gate *Si*

All basic events which do not appear in any of the minimal cut sets can be treated as fixed.

Step 2: Calculate mean and variance using gate *Si* higher-level data

$E[p_{Si}^{+(i)}]$
$\mathrm{Var}[p_{Si}^{+(i)}]$

Step 3: Begin Monte Carlo simulation

For each cycle n:

Step 3.1 Draw probability $p_{j,n}^{*(i-1)}$ for each basic event Bj (Section 2.5 describes how to generate these)
Step 3.2 Calculate probability of gate Si, $p_{Si,n}^{*(i-1)} = Si[\mathbf{p}_n^{*(i-1)}]$, where $\mathbf{p}_n^{*(i-1)}$ has elements $p_{j,n}^{*(i-1)}$
Step 3.3 Calculate $\mathrm{RIR}_{j,n}^{*(i-1)}$ for each basic event
Step 3.4 Draw probability $p_{Si,n}^{+(i)}$ from the SOK posterior distribution of the probability of gate Si based on the higher-level data
Step 3.5 Calculate back-estimate of basic event probability, $p_{j,n}^{+(i)}$ using eqn (2).

Accumulate sums over n for:

$a_{j,n}^{(i)} = [1-p_{Si,n}^{*(i-1)}\mathrm{RIR}_{j,n}^{*(i-1)}]/[1-p_{Si,n}^{*(i-1)}]$
$b_{j,n}^{(i)} = \mathrm{RIR}_{j,n}^{*(i-1)} - a_{j,n}^{(i)}$
$p_{j,n}^{+(i)} = p_{j,n}^{*(i-1)}[a_{j,n}^{(i)} + b_{j,n}^{(i)}p_{Si,n}^{+(i)}]$
$p_{j,n}^{*(i-1)}\{a_{j,n}^{(i)} + b_{j,n}^{(i)}E[p_{Si}^{+(i)}]\}$
$p_{j,n}^{+(i)}p_{j,n}^{*(i-1)}$
$(p_{j,n}^{*(i-1)}\{a_{j,n}^{(i)} + b_{j,n}^{(i)}E[p_{Si}^{+(i)}]\})^2$
$[p_{j,n}^{*(i-1)}b_{j,n}^{(i)}]^2\mathrm{Var}[p_{Si}^{+(i)}]$.

Step 4: For each basic event *Bj*, calculate and store:

$A_j^{(i)} = \mathrm{ave}[a_{j,n}^{(i)}]$
$B_j^{(i)} = \mathrm{ave}[b_{j,n}^{(i)}]$.

(We will need these for all i in Section 2.5 below.)

We can now calculate $E[p_j^{+(i)}]$ and $\mathrm{Var}[p_j^{+(i)}]$ by noting that $p_{j,n}^{+(i)}$ in eqn (2) is a function of both the random vector $\mathbf{p}_n^{*(i-1)}$ and the random variable $p_{Si}^{+(i)}$. Using the well-known conditional mean and variance formulae; namely, $E(y)=E[E(y|x)]$ and $\mathrm{Var}(y) = \mathrm{Var}[E(y|x)] + E[\mathrm{Var}(y|x)]$, we have

$$E[p_j^{+(i)}] \approx \mathrm{ave}(p_{j,n}^{*(i-1)}\{a_{j,n}^{(i)} + b_{j,n}^{(i)}E[p_{Si}^{+(i)}]\}) \quad (8)$$

and

$$\mathrm{Var}[p_j^{+(i)}] \approx \mathrm{samvar}(p_{j,n}^{*(i-1)}\{a_{j,n}^{(i)} + b_{j,n}^{(i)}E[p_{Si}^{+(i)}]\}) + \mathrm{ave}([p_{j,n}^{*(i-1)}b_{j,n}^{(i)}]^2\mathrm{Var}[p_{Si}^{+(i)}]). \quad (9)$$

Once eqn (8) has been calculated, we can then calculate $\mathrm{Cov}[p_j^{*(i-1)}, p_j^{+(i)}]$ as

$$\mathrm{Cov}[p_j^{*(i-1)}, p_j^{+(i)}] \approx \mathrm{ave}(p_{j,n}^{+(i)}p_{j,n}^{*(i-1)}) - E[p_j^{+(i)}]E[p_j^{*(i-1)}], \quad (10)$$

where $E[p_j^{*(i-1)}]$ is given by eqn (6) for gate $S(i-1)$. Note that eqns (8)–(10) are only approximations due to Monte Carlo sampling error. Using eqns (8)–(10), we can now calculate eqns (4)–(7) for gate Si.

2.4 Updated approximate mean and variance of the probability of gate *Si*

After we have calculated the means and variances of the combined basic event probabilities $p_j^{*(i)}$ in eqn (4), we can then propagate these back up the fault tree to update the mean and variance of the probability of gate Si. To do this we use the approximation formulae for the mean and variance of any function $g(x_1, x_2,\ldots, x_n)$ of n random variables $x_1, x_2,\ldots, x_n$. Expanding $g(\bullet)$ in a Taylor series about the mean values of the random variables and then truncating the series at the linear terms (see Ang and Tang [20], Section 4.3.4) produces a first-order approximation for the mean occurrence probability of gate Si:

$$E[p_{Si}^{*(i)}] \approx Si\{E[\mathbf{p}^{*(i)}]\}, \quad (11)$$

where $E[\mathbf{p}^{*(i)}]$ is a vector of expectations whose elements are given in eqn (6).

The variance of $p_{Si}^{*(i)}$ may likewise be approximated by using the corresponding variance formula. Using this formula, a first-order approximation for the variance of the occurrence probability of gate Si becomes

$$\mathrm{Var}[p_{Si}^{*(i)}] \approx \sum_j \left\{\frac{E[p_{Si}^{*(i)}][\mathrm{RIR}_j^{*(i)}-1]}{1-E[p_j^{*(i)}]}\right\}^2 \times \mathrm{Var}[p_j^{*(i)}] + \sum_j \sum_{k\neq j} \left\{\frac{E[p_{Si}^{*(i)}][\mathrm{RIR}_j^{*(i)}-1]}{1-E[p_j^{*(i)}]}\right\} \times \left\{\frac{E[p_{Si}^{*(i)}][\mathrm{RIR}_k^{*(i)}-1]}{1-E[p_k^{*(i)}]}\right\}\mathrm{Cov}[p_j^{*(i)}, p_k^{*(i)}], \quad (12)$$

where $RIR_j^{*(i)}$ in eqn (12) is evaluated at its mean $E[\mathbf{p}^{*(i)}]$ and in which we have used the fact that the required partial derivatives evaluated at the mean values are given by

$$\frac{\partial Si}{\partial p_j^{*(i)}} = Si\{E[\mathbf{p}^{*(i)}|1_j]\} - Si\{E[\mathbf{p}^{*(i)}|0_j]\}. \quad (13)$$

We will examine the performance of eqn (12) in Section 3.

2.5 Simulating from the distribution of $p^{*(i)}$

Although the above results give proper marginal means and variances for $p_j^{*(i)}$, the joint distribution of the basic event probabilities is now SOK dependent: all of the probabilities depend on the common information about the gate Si. The approach is to simulate from the dependent random vector $\mathbf{p}^{*(i)}$ (required in Step 3.1 of Section 2.3). First, sample $p_{j,n}^{*(0)}$ as if there were no higher-level data (that is, taking into account only the common correlation classes of basic events). Second, as in Step 3.4 of Section 2.3, sample higher-level data $p_{Sk,n}^{+(k)}$ for $k = 1, 2, ..., i$, from the respective SOK posterior distribution of the probability of gate Sk. Then, calculate $p_{j,n}^{*(i)}$ using the expression

$$p^{*(i)}_{j,n} = p^{*(0)}_{j,n} \prod_{k=1}^{i} \{1 - w_j^{(k)} + w_j^{(k)}[A_j^{(k)} + B_j^{(k)} p^{+(k)}_{Sk,n}]\} \quad (14)$$

where the constants $w_j^{(k)}$, $A_j^{(k)}$ and $B_j^{(k)}$, $k = 1, 2, ..., i$, are all defined and calculated in eqn (4) and Step 4 of Section 2.3. By using eqn (14), we approximately preserve the implicit dependencies between the elements of $\mathbf{p}^{*(i)}$ because eqn (14) is an approximate model for the relationship between $p_j^{*(i)}$ and the initial probabilities $p_j^{*(0)}$ using successive recursive applications of eqn (4) into which eqn (2) has been substituted. Equation (14) is an approximation because of the use of the averages $A_j^{(k)}$ and $B_j^{(k)}$ from Step 4 in Section 2.3.

3 EXAMPLE

We now consider the performance of the model in Section 2 for a simple fault tree example which was used to illustrate IRRAS fault tree solution and quantification in Appendix A of Russell *et al.* [3] It is sufficient to examine the performance of the method for the case in which higher-level data exists only for the top event in the tree; thus, $m = 1$. The reason for this is that, because of the hierarchical nature of the model in Section 2, this case represents the fundamental 'building block' for more complicated trees for which there are higher-level data at multiple gates within the tree.

A complete 2^3 factorial computer experiment is designed and used in conjunction with the example to examine the performance of the method as a function of three factors: the strength of the basic event-level data (strong or weak), the strength of the top event-level data (strong or weak), and the degree of agreement between the basic event and top event-level data (agree or disagree). The results for each of these cases are compared and used as a means of assessing the performance of the model. Thus, we consider all combinations of the three factors in the following eight cases:

Case	Basic Event-Level Data	Top Event-Level Data	Data Agree?
1	Strong	Strong	Yes
2	Strong	Strong	No
3	Weak	Weak	Yes
4	Weak	Weak	No
5	Strong	Weak	Yes
6	Strong	Weak	No
7	Weak	Strong	Yes
8	Weak	Strong	No

Figure 3 contains the example fault tree from Appendix A of Russell *et al.* [3] for which there are five minimal cut sets. If the basic events are statistically independent, it follows from the rare event approximation that the approximate probability of the top event (or top gate) is $S1(\mathbf{p}) = p1*p2 + p1*p4 + p1*p3*p5 + p2*p3*p5 + p3*p4*p5$, where $pj = P(Bj)$, $j = 1, 2, 3, 4, 5$. We further assume that our initial SOK uncertainty about each true (but unknown) value of pj is adequately modeled as a truncated lognormal random variable $p_j^{*(0)}$ with mean $E[p_j^{*(0)}]$ and 95% error factor (EF) $EF_j^{*(0)}$.

The basic event mean probabilities of occurrence for our example problem are those considered by Russell *et al.* [3]; namely, $E[p_1^{*(0)}] = 0.01$, $E[p_2^{*(0)}] = 0.02$, $E[p_3^{*(0)}] = 0.03$, $E[p_4^{*(0)}] = 0.04$, and $E[p_5^{*(0)}] = 0.05$. Unlike Russell *et al.* [3], we further assume that independent data also exist for the top event. In particular, we assume that there exists an independent source of data regarding the occurrence of the top event which may be expressed as a truncated lognormal distribution whose mean is either $S1\{E[\mathbf{p}^{*(0)}]\}$ (that is, *agreement* between the means of the higher and basic event-level data) or $30 \times S1\{E[\mathbf{p}^{*(0)}]\}$ (that is, *disagreement* between the means by a factor of 30). An EF of 2 is used here to define *strong* basic and higher event-level data, while an EF of 10 denotes *weak* basic and higher event-level data. Also note that, for simplicity, the five basic events are either all simultaneously strong or weak; mixed cases are not considered.

In each of the eight cases, the method described in Section 2 was used to calculate the combined-data mean and variance of the probability of the top event.

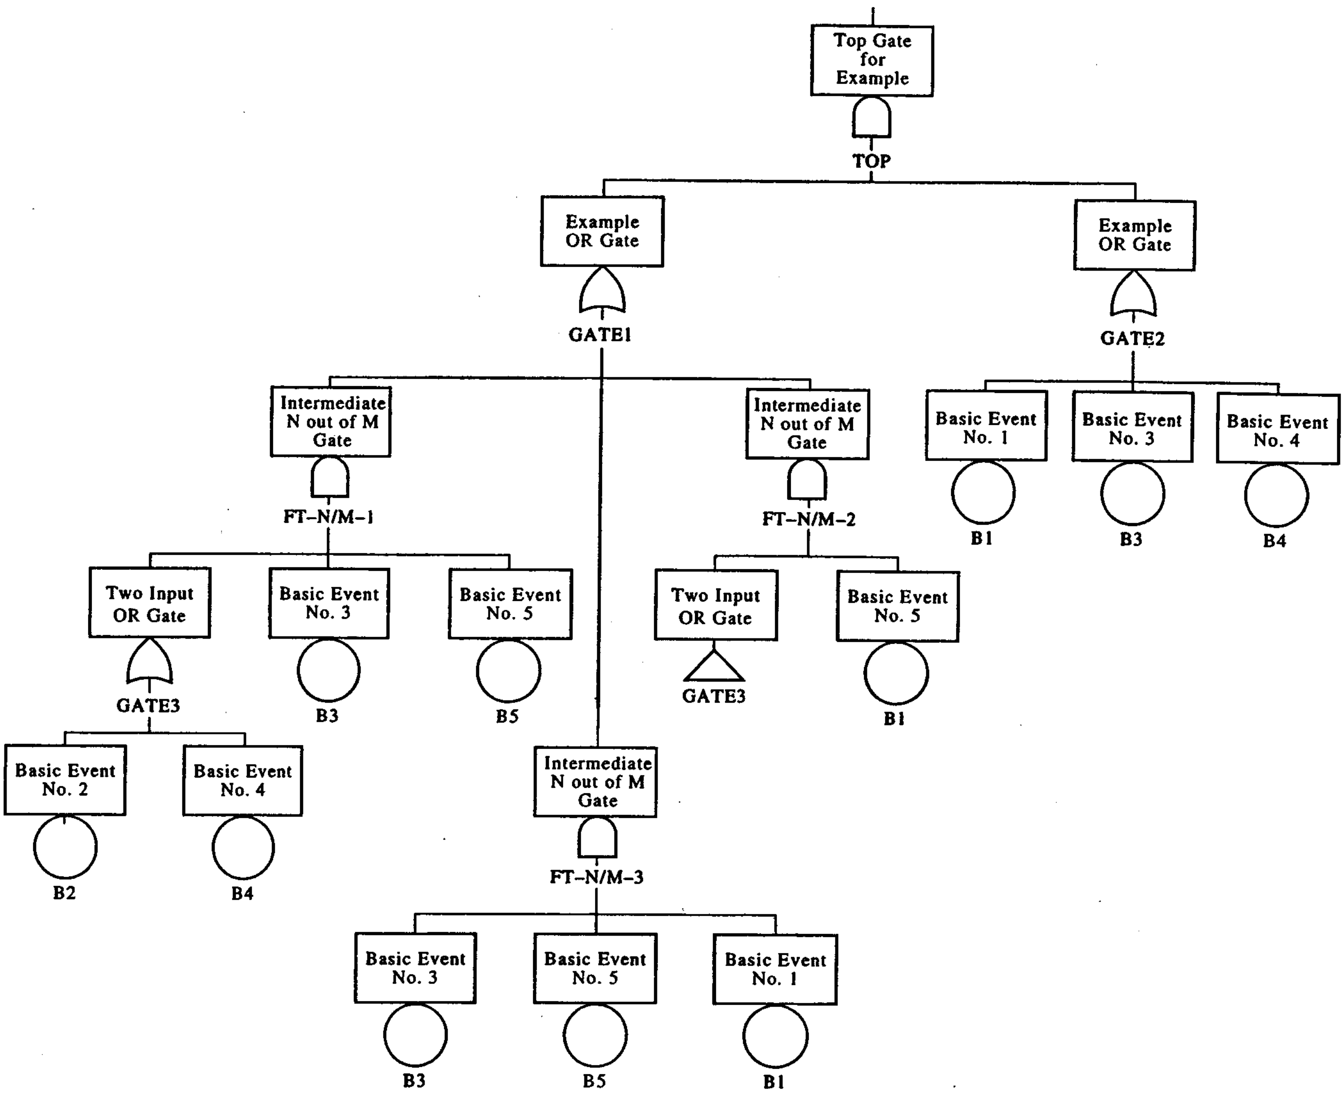

Fig. 3. Example fault tree.

In addition, two other estimates of the mean and variance of the top event probability were calculated: an estimate based *only* on the initial basic event-level data (thus ignoring the higher event-level data) and a direct estimate using *only* the top event data (thus ignoring the basic event-level data).

Sections 2.4 and 2.5 provide two methods for calculating the approximate mean and variance of the probability of the top event. Because the former method for variance calculation requires the computation of the covariances between updated basic event probabilities [eqn (12)], the latter method is easier to implement. Section 2.5 provides a method for simulating $\mathbf{p}^{*(1)}$ incorporating the SOK dependence among the basic event probabilities. Taking the sample mean and variance of $S1[\mathbf{p}^{*(1)}]$ across a Monte Carlo sample yields an estimate of the mean and variance of the probability of the top event. These estimates are subsequently referred to as the *simulated correlated* mean and variance estimates.

Section 2.4 provides first-order approximation formulae for the mean and variance of the top event probability. Note that the variance calculation formula (12) involves covariance terms. To assess the effect of ignoring this dependence, we subsequently ignore the covariance terms in eqn (12). In particular, propagating the estimates of $E\,[p_j^{*(i)}]$ and Var $[p_j^{*(i)}]$ calculated using eqns (6) and (7) (listed in Table 1) by means of eqns (11) and (12) [ignoring the covariance terms in eqn (12)] produces estimates of the mean and variance of the top event probability. These are subsequently called the *approximate* mean and variance estimates.

Table 2 contains the mean and variance of the top event probability calculated using both the simulated correlated (with a 10,000 replication Monte Carlo simulation) and the approximate estimation procedures. The results in Table 1 provide the data needed for both procedures, and Table 3 summarizes the results. The estimated means and variances are in

Table 1. Model calculation results for case 1

Basic Event	$E[p_j^{*(0)}]$	$\mathrm{Var}[p_j^{*(0)}]$	$\mathrm{Var}^{1/2}[p_j^{*(0)}]$	$p_j^{*(0)}$	$E[p_{S1}^{*(0)}]$	$\mathrm{Var}[p_{S1}^{*(0)}]$	$\mathrm{RIR}_j^{*(0)}$	$a_j^{(1)}$	$b_j^{(1)}$	$p_j^{+(1)}$	$[p_j^{*(0)} b_j^{*(1)}]^2 \mathrm{Var}[p_{S1}^{+(1)}]$	$p_j^{*(0)}\{a_j^{*(1)} + b_j^{(1)} E[p_{S1}^{+(1)}]\}$	$p_j^{+(1)} p_j^{*(0)}$
					Exact	Approximate							
B1	1.0000E − 02	1.9429E − 05	4.4078E − 03	1.0000E − 02	7.0500E − 04	1.2916E − 07	8.7362E + 01	9.3907E − 01	8.6423E + 01	9.3907E − 03	7.2124E − 08	1.0000E − 02	9.3907E − 05
B2	2.0000E − 02	7.7715E − 05	8.8156E − 03	2.0000E − 02			1.6986E + 01	9.8872E − 01	1.5997E + 01	1.9774E − 02	9.8847E − 09	2.0000E − 02	3.9549E − 04
B3	3.0000E − 02	1.7486E − 04	1.3223E − 02	3.0000E − 02	Simulated	Simulated	5.8156E + 00	9.9660E − 01	4.8190E + 00	2.9898E − 02	2.0183E − 09	3.00003 − 02	8.9694E − 04
B4	4.0000E − 02	3.1086E − 04	1.7631E − 02	4.0000E − 02	7.0406E − 04	1.3995E − 07	1.6660E + 01	9.8895E − 01	1.5671E + 01	3.9558E − 02	3.7941E − 08	4.0000E − 02	1.5823E − 03
B5	5.0000E − 02	4.8572E − 04	2.2039E − 02	5.0000E − 02			3.8298E + 00	9.9800E − 01	2.8318E + 00	4.9900E − 02	1.9359E − 09	5.0000E − 02	2.4950E − 03
$E[p_{S1}^{+(1)}]$	7.0500E − 02												
$\mathrm{Var}[p_{S1}^{+(1)}]$	9.6566E − 08												

Basic Event	$E[p_j^{+(1)}]$	$\mathrm{Var}[p_j^{+(1)}]$	$\mathrm{Cov}[p_j^{*(0)}, p_j^{+(1)}]$	$w_j^{(1)}$	$1 - w_j^{(1)}$	$E[p_j^{*(1)}]$	$\mathrm{Var}[p_j^{*(1)}]$	$RIR_j^{*(1)}$	$E[p_{S1}^{*(1)}]$	$Var[p_{S1}^{*(1)}]$	$A_j^{(1)}$	$B_j^{(1)}$
									Approximate	Approximate		
B1	9.9946E − 03	1.7689E − 05	1.8619E − 05	5.2343E − 01	4.7657E − 01	9.9972E − 03	1.8548E − 05	8.7391E + 01	7.0427E − 04	1.2599E − 07	9.3916E − 01	9.9942E + 01
B2	1.9963E − 02	7.5242E − 05	7.4665E − 05	5.0808E − 01	4.9192E − 01	1.9981E − 02	7.5552E − 05	1.6997E + 01			9.8872E − 01	1.7698E + 01
B3	3.0120E − 02	1.7858E − 04	1.8211E − 04	4.9474E − 01	5.0526E − 01	3.0060E − 02	1.7940E − 04	5.8045E + 00			9.9663E − 01	5.3150E + 00
B4	3.9953E − 02	3.1604E − 04	3.1670E − 04	4.9587E − 01	5.0413E − 01	3.9977E − 02	3.1505E − 04	1.6671E + 01			9.8895E − 01	1.7394E + 01
B5	4.9736E − 02	4.8776E − 04	4.7442E − 4	4.9895E − 01	5.0105E − 01	4.9868E − 02	4.8058E − 04	3.8369E + 00			9.9799E − 01	3.1577E + 00

Table 2. Monte Carlo simulation results for the mean and variance of the top event probability for case 1

Basic Event	$E[p_j^{*(1)}]$	$\mathrm{Var}[p_j^{*(1)}]$	$w_j^{(1)}$	$1-w_j^{(1)}$	$A_j^{(1)}$	$B_j^{(1)}$	$p_j^{*(0)}$	$p_j^{*(1)}$	$E[p_{SI}^{*(1)}]$	$\mathrm{Var}[p_{SI}^{*(1)}]$
B1	9.9972E − 03	1.8548E − 05	5.2343E − 01	4.7657E − 01	9.3916E − 01	9.9942E + 01	1.0000E − 02	1.0050E − 02		
B2	1.9981E − 02	7.5552E − 05	5.0808E − 01	4.9192E − 01	9.8872E − 01	1.7698E + 01	2.0000E − 02	2.0012E − 02	Simulated	Simulated
B3	3.0060E − 02	1.7940E − 04	4.9474E − 01	5.0526E − 01	9.9663E − 01	5.3150E + 00	3.0000E − 02	3.0006E − 02	Correlated	Correlated
B4	3.9977E − 02	3.1505E − 04	4.9587E − 01	5.0413E − 01	9.8895E − 01	1.7394E + 01	4.0000E − 02	4.0024E − 02	7.1316E − 04	1.4204E − 07
B5	4.9868E − 02	4.8058E − 04	4.9895E − 01	5.0105E − 01	9.9799E − 01	3.1577E + 00	5.0000E − 02	5.0005E − 02		
$E[p_{SI}^{+(1)}]$	7.050 0E − 04								Approximate	Approximate
$\mathrm{Var}[p_{SI}^{+(1)}]$	9.6566E − 08								7.0427E − 04	1.2599E − 07

Table 3. Summary of results for case 1

Basic Event	$E[p_j^{*(0)}]$	$\mathrm{Var}[p_j^{*(0)}]$	$E[p_j^{+(1)}]$	$\mathrm{Var}[p_j^{+(1)}]$	$E[p_j^{*(1)}]$	$\mathrm{Var}[p_j^{*(1)}]$	$E[p_{SI}^{*(0)}]$	$\mathrm{Var}[p_{SI}^{*(0)}]$	$E[p_{SI}+^{(1)}]$	$\mathrm{Var}[p_{SI}^{+(1)}]$	$E[p_{SI}^{*(1)}]$	$\mathrm{Var}[p_{SI}^{*(1)}]$
							Exact	Approximate	Exact	Exact	Simulated	Simulated
B1	1.0000E − 02	1.9429E − 05	2.6680E − 02	5.6644E − 04	9.9972E − 03	1.8548E − 05	7.0500E − 04	1.2916E − 07	7.0500E − 04	9.6566E − 08	Correlated	Correlated
B2	2.0000E − 02	7.7715E − 05	2.7226E − 02	1.7121E − 03	1.9981E − 02	7.5552E − 05					7.1316E − 04	1.4204E − 07
B3	3.0000E − 02	1.7486E − 04	3.2565E − 02	3.6155E − 03	3.0060E − 02	1.7940E − 04	Simulated	Simulated				
B4	4.0000E − 02	3.1086E − 04	4.9274E − 02	5.6066E − 03	3.9977E − 02	3.1505E − 04	7.0406E − 04	1.3995E − 07			Approximate	Approximate
B5	5.0000E − 02	4.8572E − 04	4.9576E − 02	7.6368E − 03	4.9868E − 02	4.8058E − 04					7.0427E − 04	1.2599E − 07

good agreement despite the fact that we have ignored the covariances in eqn (12). Although in case 1 the dependencies induced between the elements of $\mathbf{p}^{*(1)}$ due to the common use of the top event-level data are sufficiently small so that they can essentially be ignored, this is not always the case, as we shall see below.

Figure 4 shows the mean top event probability for all eight cases using all three estimation methods. For convenience, the values of the three factors for each case are also shown at the bottom of Fig. 4. The results for the combined data in Fig. 4 are all based on the 'simulated correlated' estimates described above and as illustrated in Table 2. The length of the 'whisker' extending above each bar in Fig. 4 represents the standard deviation associated with the corresponding estimate; for example, $\mathrm{Var}^{1/2}[p_{Si}^{*(1)}]$ associated with $E\,[p_{Si}^{*(1)}]$. Similarly, Fig. 5 contains the corresponding results based on using the 'approximate' mean and variance estimates described above for the combined data.

Upon examining and comparing Figs 4 and 5, several things become apparent. First, consider Fig. 4. If the initial basic event data are strong, then the addition of top event data has little effect on the variance of the top event probability (cases 1, 2, 5 and 6). On the other hand, if the initial basic event data are weak, then the additional use of top event data can have a significant effect on the variance of the top event probability (cases 3, 4, 7 and 8), particularly so when there is no agreement between the basic event and top event data (cases 4 and 8). If one source is weak while the other is strong and they don't agree, then the mean top event probability is an average of the corresponding means for the two sources and tends to be rather dramatically pulled towards the mean of the stronger source (cases 6 and 8). The case where both sources are strong but both disagree should not be combined before first rectifying the disagreement; otherwise, the variance of the top event probability may be underestimated (case 2).

Now compare Figs 4 and 5. The mean top event probabilities based on the combined data disagree in Figs 4 and 5 when the initial basic event data are weak

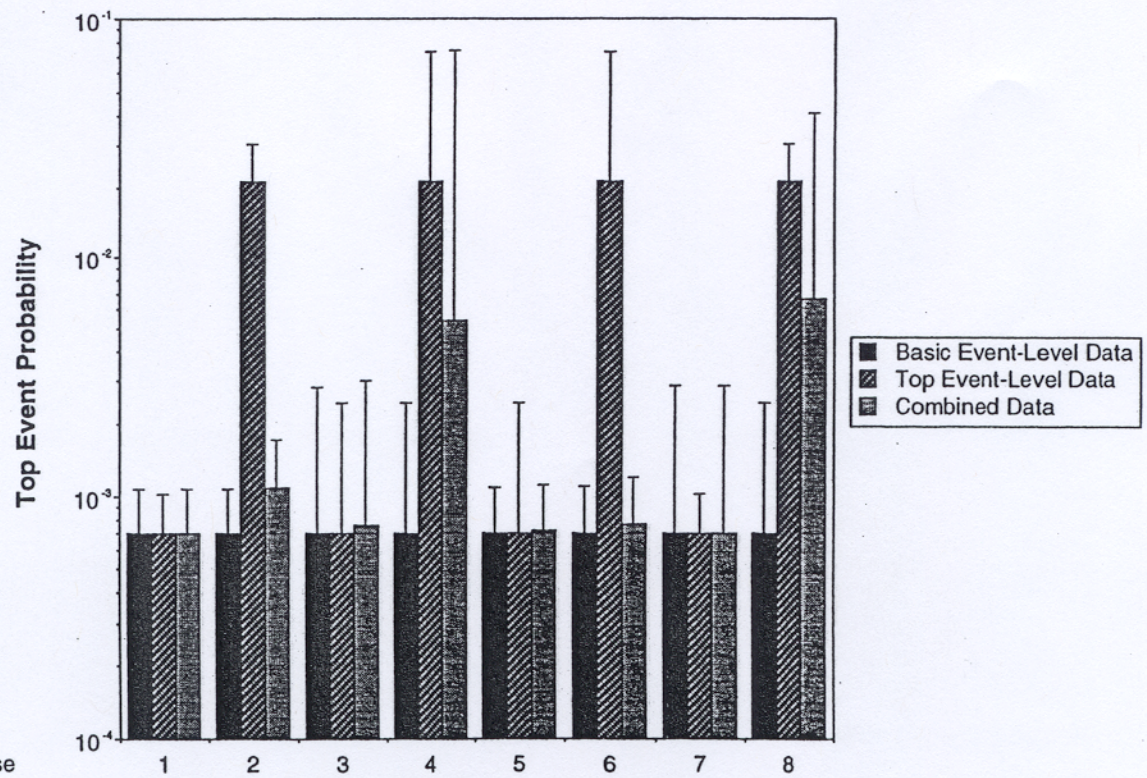

Case	1	2	3	4	5	6	7	8
Basic Event-Level Data	Strong	Strong	Weak	Weak	Strong	Strong	Weak	Weak
Top Event-Level Data	Strong	Strong	Weak	Weak	Weak	Weak	Strong	Strong
Data Agree?	Yes	No	Yes	No	Yes	No	Yes	No

Fig. 4. A comparison of the mean top event probability using different data sources for eight cases based on correlated Monte Carlo simulation for the combined data.

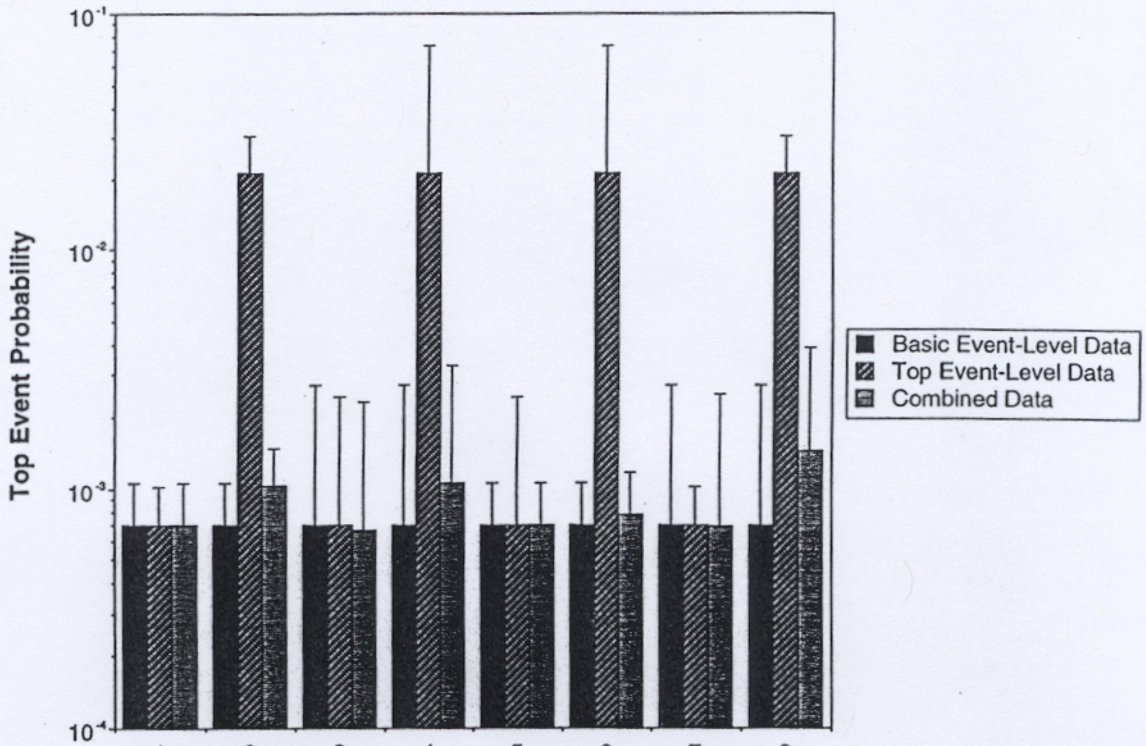

Case	1	2	3	4	5	6	7	8
Basic Event-Level Data	Strong	Strong	Weak	Weak	Strong	Strong	Weak	Weak
Top Event-Level Data	Strong	Strong	Weak	Weak	Weak	Weak	Strong	Strong
Data Agree?	Yes	No	Yes	No	Yes	No	Yes	No

Fig. 5. A comparsion of the mean top event probability using different data sources for eight cases based on mean and variance approximation formulae for the combined data.

and both data sources disagree (compare cases 4 and 8 in Figs 4 and 5). In this case, we see that the dependencies induced between the elements of $\mathbf{p}^{*(1)}$ are such that ignoring them tends to underestimate the mean top event probability. The variance approximation formula (12), in which the covariances are ignored, likewise underestimates the variance of the top event probability in almost all cases, especially when the initial basic event data are weak (compare cases 3, 4, 7 and 8 in Figs 4 and 5), and severely so when the initial basic event data are weak and both sources disagree (compare cases 4 and 8 in Figs 4 and 5). Finally, the mean and variance approximation formulae, in which the covariances are ignored, give reasonably accurate results as long as the initial basic event data are strong relative to the top event-level data (compare cases 1, 2, 5 and 6 of Figs 4 and 5).

4 CONCLUSIONS

A top-down methodology has been developed for using higher-level failure data in fault tree quantification. The method requires the identification and use of SOK uncertainty distributions for the probabilities of occurrence of both the initial basic and higher-level events. Such identification and use is consistent with the SOK distribution assumptions required in uncertainty analysis. The top-down structure of the model makes it relatively easy to implement in existing fault tree quantification and uncertainty analysis codes, such as IRRAS [4, 5].

The performance of the model is illustrated for a simple example and performs as expected. The combined use of higher-level data is particularly advantageous when the initial basic event data are weak. When the initial basic event-level data disagree with the higher event-level data, the combined probability estimate is essentially a weighted average of the probability estimated from each data source with weights proportional to the strength of each data source.

The model involves updating and propagating only the first two moments of SOK uncertainty distributions on event probabilities. A future task would be to extend this to include updating the SOK distributions themselves. Although this is a complicated analytical task, it can be numerically accomplished using Markov chain simulation, also known as Markov chain Monte Carlo [21].

Although we have exclusively considered fault tree models, the method is also directly applicable to networks, trees, and similar graphical models [22, 23]. In particular, the method can be used in the tree of cliques propagation algorithm described by Almond [22].

A major limitation of the method described in this paper is that it only works when the higher-level data are independent of the basic event-level probabilities. In particular, basic and higher event-level data from the same series of tests or common exposure time cannot both be simultaneously used to update the model. On the other hand, when test or experiential data are available simultaneously at both levels, it provides us with an opportunity to validate the structure of the model, particularly the completeness and adequacy of the fault tree.

ACKNOWLEDGEMENT

This work was supported by the US Nuclear Regulatory Commission, Office of Nuclear Regulatory Research, under Job Code Number W6489-5 (Dr Lee R. Abramson, NRC Technical Monitor) at Los Alamos National Laboratory. We are grateful for this support.

We are also grateful to two referees, whose thought provoking and helpful comments improved the presentation.

REFERENCES

1. Vesely, W. E., Goldberg, F. F., Roberts, N. H. and Haasl, D. F., Fault tree handbook. NUREG-0492, US Nuclear Regulatory Commission, Washington DC, 1981.
2. Hickman, J. W., *PRA procedures guide: A guide to the performance of probabilistic risk assessments for nuclear power plants*, Vol. 1. American Nuclear Society La Grange Park, Illinois, and Institute of Electrical and Electronic Engineers, New York, NY, NUREG/CR-2300, 1983.
3. Russell, K. D., Atwood, C. L., Galyean, W. J., Sattison, M. B. and Rasmuson, D. M., Systems analysis programs for hands-on integrated reliability evaluations (SAPHIRE) version 5.0, vol. 1-technical reference manual. EGG-2716 (NUREG/CR-6116), Idaho National Engineering Laboratory, 1994.
4. Russell, K. D., Kvarfordt, K. J., Skinner, N. L., Wood, S. T. and Rasmuson, D. M., Systems analysis programs for hands-on integrated reliability evaluations (SAPHIRE) version 5.0, vol. 2-Integrated Reliability and Risk Analysis System (IRRAS) reference manual. EGG-2716 (NUREG/CR-6116), Idaho National Engineering Laboratory, 1994.
5. VanHorn, R. L., Russell, K. D. and Skinner, N. L., Systems analysis programs for hands-on integrated reliability evaluations (SAPHIRE) version 5.0, vol. 3-Integrated Reliability and Risk Analysis System (IRRAS) tutorial manual. EGG-2716 (NUREG/CR-6116), Idaho National Engineering Laboratory, 1994.
6. Grant, G. M. Roesener, W. S., Hall, D. G., Atwood, C. L., Gentillon, C. D. and Wolf, T. R., High-pressure coolant injection (HPCI) system performance, 1987-1993. INEL-94/0158, Idaho National Engineering Laboratory, 1995.
7. Easterling, R. G. & Prairie, R. R., Combining component and system information. *Technometrics*, 1971, **13**, 271–279.
8. Mastran, D. V., Incorporating component and system test data into the same assessment: a Bayesian approach. *Operations Research*, 1976, **24**, 491–499.
9. Mastran, D. V. & Singpurwalla, N. D., A Bayesian estimation of the reliability of coherent structures. *Operations Research*, 1978, **26**, 663–672.
10. Barlow, R. E., Combining component and system information in system reliability calculation. In *Probabilistic Methods in the Mechanics of Solids and Structures*, eds S. Eggwertz and N. L. Lind. Springer-Verlag, Berlin, 1985, pp.375–383.
11. Martz, H. F., Waller, R. A. & Fickas, E. T., Bayesian reliability analysis of series systems of binomial subsystems and components. *Technometrics*, 1988, **30**, 143–154.
12. Martz, H. F. & Waller, R. A. Bayesian reliability analysis of complex series/parallel systems of binomial subsystems and components. *Technometrics*, 1990, **32**, 407–416.
13. Martz, H. F. and Baggerly, K. A., Bayesian reliability analysis of high-reliability systems of binomial and Poisson subsystems. LA-UR-95-2855, Los Alamos National Laboratory, 1995.
14. Hulting, F. L. & Robinson, J. A., The reliability of a series system of repairable subsystems: a Bayesian approach. *Naval Research Logistics*, 1994, **41**, 483–506.
15. Bier, V. M., On the concept of perfect aggregation in Bayesian estimation. *Reliability Engineering and System Safety*, 1994, **46**, 271–281.
16. Azaiez, M. N. & Bier, V. M., Perfect aggregation in Bayesian estimation for several realistic reliability models. *Applied Mathematics and Computation*, 1994, **64**, 139–153.
17. Azaiez, M. N. & Bier, V. M., Perfect aggregation for a class of general reliability models with Bayesian updating. *Applied Mathematics and Computation*, 1995, **73**, 281–302.
18. Haim, M., The impact of new evidence on state-of-knowledge dependence. *Reliability Engineering and System Safety*, 1993, **40**, 1–4.
19. Apostolakis, G. & Moieni, P., The foundations of models of dependence in probabilistic safety assessment. *Reliability Engineering*, 1987, **18**, 177–194.
20. Ang, A. and Tang, W., *Probability Concepts in Engineering Planning and Design* (*Vol.* 1 *Basic Principles*). John Wiley, New York, 1975.
21. Gilks, W. R., Richardson, S. and Spiegelhalter, D. J., *Markov Chain Monte Carlo in Practice*. Chapman and Hall, London, 1996.
22. Almond, R. G., *Graphical Belief Modeling*. Chapman and Hall, London, 1995.
23. Jensen, F. V., *An Introduction to Bayesian Networks*. Springer, New York, 1996.