Selection of risk reduction portfolios under interval-valued probabilities

Antti Toppila and Ahti Salo

Systems Analysis Laboratory, Department of Mathematics and Systems Analysis, Aalto University School of Science, P.O. Box 11100, 00076 Aalto, FINLAND

Abstract

A central problem in risk management is to identify the optimal combination (or portfolio) of actions that improves the reliability of the system most through reducing failure event probabilities, subject to the availability of resources. This optimal portfolio can be sensitive with regard to epistemic uncertainties about the failure events' probabilities. In this paper, we develop an optimization model to support the allocation of resources to risk mitigation actions in coherent systems in which interval-valued probabilities defined by lower and upper bounds are employed to capture epistemic uncertainties. Decision recommendations are based on portfolio dominance: a resource allocation portfolio is dominated if there exists another portfolio that improves system reliability (i) at least as much for all feasible failure probabilities and (ii) strictly more for some feasible probabilities. Based on non-dominated portfolios, recommendations about actions to implement are derived by inspecting in how many non-dominated portfolios a given improvement is contained. We present an exact method for computing the non-dominated portfolios. We also present an approximate method that simplifies the reliability function using total order interactions so that larger problem instances can be solved with reasonable computational effort.

Keywords: epistemic uncertainty, interval-valued probabilities, reliability allocation

1. Introduction

The reliability of a system can be improved by using higher quality components, introducing redundant components, reducing the operational and environmental loads on the components, or improving the maintainability of the component, for instance (Rausand and Høyland, 2004). The selection of components to be improved can be formalized as the *redundancy allocation problem*: Find which combination (or *portfolio*) of improvements increases the reliability the most, given that only some of the possible improvements can be implemented subject to of relevant economical, technical or other constraints (such as the available budget for improvements). In risk informed decision making, this problem guides the allocation of resources to improvement activities so that the reliability of the system improves as much as possible.

In this paper, we consider the redundancy allocation problem under epistemic uncertainty about probabilities (Apostolakis, 1990). This type of uncertainty – which is present in most redundancy allocation models – is also referred to as reducible, or non-stochastic uncertainty, which expresses the lack of knowledge about the probabilities experienced by the decision-maker or the analyst (see also Walley, 1991). In reliability analysis, epistemic uncertainty has been modeled with methods such as Dempster-Shafer theory and imprecise probability/reliability (Limbourg et al., 2008; Sallak et al., 2013; Utkin and Coolen, 2007a; Helton et al., 2004; Coolen, 2004; Le Duy et al., 2010, 2011).

Risk importance measures such as the Fussell-Vesely and Birnbaum measures have been used for supporting the selection of improvement actions under interval-valued probabilities (Toppila and Salo, 2013; Kuo and Zhu, 2012b). These measures can be used for guiding the selection of an improvement portfolio so that improvements are added to the portfolio in decreasing importance starting from the most important, then adding the second most important, and so on, until no more improvements can be added to the portfolio without violating resource or other constraints. This so called greedy heuristic can be used for finding a portfolio that yields a good (although not optimal) portfolio.

However, the drawback with the greedy heuristic is that an improvement has a non-linear impact on the reliability of the system. For instance, consider a series system of parallel systems, where each parallel system consists of k identical and independent components, and the components are 50 % reliable without improvement and 100 % reliable with the improvement. By symmetry, every component has the same importance, which would suggest that the components that are improved do not matter, only how many of them are improved. Clearly, this conclusion is incorrect, because by improving a single component from every parallel system, the entire system can be made 100 % reliable, whereas leaving one parallel system without any improvements leads to a system that has reliability less than 100 %, even if components in the remaining parallel systems are improved to 100 % reliable.

In this paper, we develop a method for solving the redundancy allocation problem under interval valued probabilities. This method is based on the implicit enumeration of all possible portfolios, and we use it to determine the optimal redundancy allocation under uncertainty about the probabilities. The uncertainty is characterized with lower and upper bounds that determine an interval called the *feasible set* to which the probability belongs. The values that the system reliability can attain when the probabilities are within their respective feasible sets define a corresponding interval. Our approach resembles the Robust Portfolio Modeling method (RPM; Liesiö et al., 2007; Liesiö et al., 2008) which offers robust decision recommend ations for multicriteria project selection based on incomplete information about criterion weights and interval-valued statements about projects' performance on these criteria. In optimal reliability allocation, analogous problems have been considered by Feizollahi and Modarres (2012) who optimize the reliability of a series system of parallel structures under interval uncertainty about parameter values.

We present an exponential time solution algorithm for this problem and solve an illustrative case that represents a residual heat removal system of a nuclear power plant. This example has 31 failure events, the probabilities of which could be reduced through risk management actions. When the probabilities' intervals

were very wide, our algorithm did not solve all non-dominated portfolios within reasonable time, when considering portfolios of more than five simultaneous improvements. For such situations, we use total order interactions introduced by Borgonovo and Smith (2011) and Borgonovo (2010) to simplify the reliability function of the system, and compute the non-dominated (ND) set with respect to it.

The rest of this paper is structured as follows. Section 2 reviews methods for redundancy allocation. Section 3 formulates the optimal redundancy allocation problem as a portfolio optimization problem. Section 4 covers computational experiments related to optimal redundancy allocation and the analysis of epistemic uncertainties. Section 5 presents a method for simplifying the reliability function so that larger problem instances can be solved. Section 6 discusses and Section 7 concludes.

2. Exact methods for optimal redundancy allocation

Redundancy and reliability allocation has attracted a significant amount of research since 1960's (for reviews, see Kuo and Wan, 2007; Kuo and Prasad, 2000; Tillman et al., 1977). This research has resulted in different kinds of optimization methods ranging from heuristic algorithms (e.g. genetic algorithms, ant colony optimization) to exact algorithms.

In this paper, we consider exact methods for optimal redundancy allocation; these include dynamic programming, implicit enumeration algorithms and branch-and-bound algorithms. The advantage of these methods is that they yield a globally optimal solution. On the downside, they are usually computationally more demanding than heuristic methods, which may limit their usefulness in the analysis of large systems. In general, the redundancy allocation problem of a series system of parallel structures has been proven to be NP-hard (Chern, 1992), and thus the worst case solution time of any algorithm for this problem increases exponentially as a function of problem size (assuming the likely case that the famous conjecture $P \neq NP$ is valid).

The solution approaches can be divided into methods that assume a specific structure of the system and those that do not. Specific structures are for instances series-parallel (e.g. Sung and Cho, 1999) and k out of n systems (e.g. Elegbede et al., 2003), which are common in many applications of reliability optimization. Generic optimization methods that can solve complex systems are typically formulated as integer non-linear and non-convex programming problems (e.g. Kuo and Wan, 2007). Methods for solving them are tailored versions of generic exact algorithms such as branch-and-bound and implicit enumeration.

Prasad and Kuo (2000) present a method for optimizing the reliability of coherent systems with statistically independent component failures. Their method is based on implicit enumeration of all alternatives in lexicographical order. They exploit the property that the reliability of a coherent system increases when there is more redundancy. Specifically, they formulate a bounding rule which can eliminate parts of the solution space without the need to evaluate these parts explicitly. They apply their method to a parallel system, a series system and more complex (coherent) system under multiple constraints that are monotone with respect to the number of redundant components.

Ha and Kuo (2006) develop a branch-and-bound algorithm to solve the optimal redundancy allocation problem in a coherent system in which improvements are formed by adding redundant components to the system under multiple constraints on how many redundant components can be used. These constraints are assumed to be monotone with respect to the number of redundant components: that is, if all variables are on the left-hand-side of the constraint equation, then the function defined by the left-hand-side is a monotone function. They report a significant increase in optimization speed compared to the method of Prasad and Kuo (2000). The speed-up is attained through an efficient bounding scheme, which is based on a greedy heuristic that efficiently eliminates branches from the search tree. They also note that their method could be further improved for problem instances in which the structure of the system can be exploited to make the bounding process more efficient.

3. Methodological development for solving the optimal risk reduction portfolio problem

Consider a system with the reliability function $r(p)$, where $p = (p_1, \ldots, p_n)$ such that p_i is the probability of event $i = 1, \ldots, n$. As proven by Borgonovo (2010), $r(p)$ is multilinear with respect to its arguments for both coherent and non-coherent systems, and thus

$$r(p) = \sum_{I \subseteq \{1,\ldots,n\}} \alpha_I \prod_{j \in I} p_j \tag{1}$$

for some coefficients $\alpha_I \in \mathbb{R}$, where I is a set valued index. In typical applications, the multilinear function is sparse so that most of the coefficients α_I are zero. Thus from computational perspective, it is preferable to use a sparse data structure that saves only the non-zero coefficients and the indices of the variables that are included in the corresponding term. For instance, in many programming languages, $r(p) = 1 - p_1 p_2$ can be represented by the coefficient vector $(1, -1)$ and the corresponding list of lists $((), (1, 2))$, where $()$ denotes the empty list, and the vector and list of lists can be easily implementable by built-in commands.

Assume that the system can be altered by improvements which each reduce the probability of a single failure event. The improvement i changes the probability of the associated failure event from p_i to p_i^r for a cost c_i. We encode a portfolio $P \subseteq \{1, \ldots, n\}$ of improvements as the vector $z = (z_1, \ldots, z_n)$ such that

$$z_i = \begin{cases} 1 & \text{if } i \in P \\ 0 & \text{if } i \notin P. \end{cases}$$

Because there is a one-to-one mapping between the binary vector z and the corresponding set P, we also call z a portfolio. Under this notation, the set of all portfolios can be expressed as $\mathcal{Z} = \{0, 1\}^n$. The cost of a portfolio is $\sum_i c_i z_i$. A portfolio is feasible if its cost does not exceed a given budget B.

The goal is to determine the portfolio z that reduces the risk – or equivalently increases the reliability of the system – as much as possible. More specifically,

we define

$$r(z;p) = r(q) \ ,$$

where $q_i = p_i(1 - z_i) + p_i^r z_i$. Then for a given p, the optimal portfolio is the solution to the problem

$$\begin{aligned} \max_{z} \quad & r(z;p) \\ \text{s.t.} \quad & \sum_{i=1}^{m} c_i z_i \leq B \ . \end{aligned} \tag{2}$$

This is a knapsack optimization problem with a non-linear objective function and a linear budget constraint.

3.1. Interval-valued probabilities and dominance

As noted in Section 1, it is of interest to analyze how epistemic uncertainty about the probabilities captured by interval-values impacts the results of the analysis. Thus, assume that for $i = 1, \ldots, n$ that $p_i \in [\underline{p}_i, \overline{p}_i]$, or equivalently $\underline{p} \leq p \leq \overline{p}$, where $\underline{p}_i$ and $\overline{p}_i$ are the lower and upper bounds of p_i. The set of feasible probabilities is $\mathcal{P}_F = \{p \in \mathbb{R}^m | \underline{p}_i \leq p_i \leq \overline{p}_i\}$. These bounds can be inferred, for instance, from the confidence intervals of the events' probability estimates, simulation (Borgonovo, 2008; Modarres, 2006; Zio, 2011; Baraldi et al., 2009; Modarres, 2006), interval-probabilities (Weichselberger, 2000), fuzzy probabilities Buckley (2003), coherent lower and upper probabilities (Walley, 1991), imprecise reliability (Utkin and Coolen, 2007b), and Dempster-Shafer theory (Dempster, 1967; Shafer, 1976), depending on what type of models are preferred by the analyst.

In our model with interval-valued probabilities, we focus on consistent improvements. For instance, with wide intervals, if p_i and p_i^r have overlapping intervals within which the probabilities can vary freely, then the probability of an event with the improvement may not actually be smaller than what it was before the improvement. In our computational model, we assume that the system is monotone with respect to improvements, wherefore the impact of the improvements such that including event i to the improvement portfolio reduces the probability of event i by a factor of $a_i, 0 \leq a_i \leq 1$, that is the probability of

the event after the change is $p_i^r = a_i p_i \in [a_i \underline{p}_i, a_i \overline{p}_i]$. Here, $a_i = 0$ represents a change that makes event i impossible, whereas $a_i = 1$ means that the probability of event i does not change. If an improvement will increase the probability of the event, then this event is replaced by the complement of event i, and the technique can be applied to that event, resulting in an equivalent analysis.

With interval-valued unreliabilities, the solution to problem (2) is no longer unique, because with different values of p, the optimal portfolio may be different. However, some portfolios are inferior to others in that they are less reliable for *all* feasible p. For such settings, we define dominance as follows.

Definition 1. Portfolio z dominates portfolio z', denoted $z \succ z'$, iff $r(z;p) \geq r(z;p')$ for all $\underline{p} \leq p \leq \overline{p}$, and $r(z;p) > r(z';p))$ for some $\underline{p} \leq p \leq \overline{p}$. The set of non-dominated portfolios is

$$\mathcal{Z}_N = \{z \in \mathcal{Z}_F |\ \nexists z' \in \mathcal{Z}_F : z' \succ z\} \ ,$$

where $\mathcal{Z}_F = \{z \in \mathcal{Z} | \sum_{i=1}^m c_i z_i \leq B\}$ is the set of feasible portfolios.

A rational decision maker can limit her choices to non-dominated portfolios (see, e.g. Liesiö et al., 2007; Liesiö et al., 2008; Toppila and Salo, 2013). This is because if she were to select a dominated portfolio, then it would be possible to select some non-dominated portfolio of improvements such that the system is as reliable for all feasible probabilities and more reliable for some feasible probabilities than this selected portfolio.

If the information about the probabilities becomes more complete so that the lower bound of the probability is increased and/or upper bound is decreased (resulting in a narrower interval), then any portfolio that is dominated will stay dominated, and the new set of non-dominated portfolios will be a subset of the current non-dominated set. Thus, if these intervals do contain the values of the probabilities, then the optimal portfolio corresponding to these estimates is within the set of non-dominated portfolios.

The set of non-dominated portfolios can provide guidance on which improvements should be included in a portfolio through the use of core indices(Liesiö et al., 2007).

Definition 2. The core index CI of improvement i is defined as

$$CI_i = \frac{|\{z \in \mathcal{Z}_N \mid z_i = 1\}|}{|\mathcal{Z}_N|},$$

where $|.|$ denotes the number of elements in a set.

The core indices divide improvements into three categories, namely core ($CI_i = 1$), exterior ($CI_i = 0$) and borderline improvements ($0 < CI_i < 1$). Core improvements are in all non-dominated portfolios, wherefore any rational decision maker will choose these improvements in her portfolio (because otherwise the selected portfolio cannot be non-dominated). On the other hand, a rational decision maker will not select an exterior improvement into her portfolio, because doing so would also make her portfolio dominated (because there does not exist a non-dominated portfolio that would include these improvements). Borderline improvements could be included or excluded from the portfolio depending on what other improvements are in the portfolio without necessarily resulting in a dominated or non-dominated portfolio. These are also the most interesting improvements for elicitation of tighter bounds on the probabilities, because these are the improvements which selection is uncertain.

In the context of the RPM methodology, several suggestions for visualizing and interactively exploring the set of non-dominated portfolios have been presented. This development is straightforward once the non-dominated set has been computed (Liesiö et al., 2007; Liesiö et al., 2008).

3.2. Computation of dominance

In our computational development, we build upon the framework by Toppila and Salo (2013) in which dominance is evaluated by examining the non-negativity of a multilinear function. Specifically, in order to check if portfolio z dominates portfolio z', the first condition in Definition 1 is $r(z,p) \geq r(z',p)\ \forall p \in \mathcal{P}_F$, which is equivalent to the non-negativity of the multilinear function $g(p) = r(z,p) - r(z',p)$. The second condition of dominance follows from the first condition if all intervals are non-singleton sets (Proposition 1 in Toppila and Salo, 2013). If some of the intervals are singleton sets (i.e. point

estimates), then these probabilities would have unique values that can be substituted into $g(p)$. This substitution results in a new multilinear function g', which contains only variables that are constrained in non-singleton sets. Thus dominance can be characterized by the non-negativity of a multilinear function.

The non-negativity of a multilinear function can be checked with the algorithm by Toppila and Salo (2013). This recursive algorithm is based on the recognition that the extreme values of a multilinear function in a hyperrectangle are obtained at the extreme points of the hyperrectangle. One recursion involves a bounding phase and if bounding is unsuccessful, a branching phase. In the bounding phase, it is checked if an easily computed lower bound is greater than zero or an easily computed least upper bound is less than zero (easily meaning computationally inexpensive). In the first case, dominance holds while in the latter case it does not. If neither of these cases holds, then a variable is branched: A variable is selected, set at its lower and upper bound, and the two new multilinear functions resulting from these substitutions are recursively tested with the test above. Because at every recursion step, a variable is eliminated, the algorithm is guaranteed to terminate with the correct answer in a finite number of iterations. Toppila and Salo (2013) also describe a method based on checking the non-negativity of the multilinear coefficient of the branching variable, which can reduce the number of branches needed.

In this paper we adapt the algorithm Toppila and Salo (2013) as follows. In this algorithm (on line 4), the (bound of the) least upper bound is computed as follows: The terms of $g(p)$ are arranged such that $g(p) = a_\emptyset + g^+(p) + g^-(p)$, where $g^+(p)$, and $g^-(p)$ denote the sums of product terms of $g(p)$ with positive, and negative coefficients α_J, respectively. An upper bound for this function is then computed as $a_\emptyset + g^+(\underline{p}) + g^-(\overline{p})$, where $\underline{p} = (\underline{p}_1, \ldots, \underline{p}_n)$, and $\overline{p} = (\overline{p}_1, \ldots, \overline{p}_n)$. However, this bound is weak if the function is separable. For instance, the bound computed on $g(p) = p_1 - p_2$, when $\underline{p}_i = 0$ and $\overline{p}_i = 1, i = 1, 2$ would yield a least bound of 0, which would not allow the algorithm to terminate without branching. To overcome this, we use the improved upper

bound computed as follows: Consider the vectors p^+ and p^-, where

$$p_i^+ = \begin{cases} \underline{p}_i, \text{if } p_i \text{ is only in terms with positive coefficient } \alpha_I \\ \overline{p}_i, \text{otherwise} , \end{cases}$$

and

$$p_i^- = \begin{cases} \overline{p}_i, \text{if } p_i \text{ is only in terms with negative coefficient } \alpha_I \\ \underline{p}_i, \text{otherwise} . \end{cases}$$

Then the improved bound that replaced on line 4 of the original algorithm is given by $\alpha_\emptyset + g^+(p^+) + g^-(p^-)$. If all variables are in terms with both positive and negative coefficients, then this bound is the same as in Toppila and Salo (2013), because then $p^+ = \overline{p}$ and $p^- = \underline{p}$. Otherwise, a variable that is only in terms with positive(negative) coefficient will affect the value of $g^+(p)(g^-(p))$ only, so that the least upper bound will be obtained when this variable is on its lower(upper) bound, which proves the correctness of the bound.

3.3. Computation of non-dominated portfolios

The set of all non-dominated portfolios can be computed by enumerating all the possible portfolios and by performing pairwise dominance checks. However, the number of portfolios increases exponentially with the number of possible improvements, and thus this approach requires too much computational resources. For instance, n improvements can form 2^n portfolios and pairwise comparisons between these can require up to $2^{2n} - 2^n$ comparisons (for $n = 4$ this yields 240 comparisons and $n = 10$ about one million comparisons). This problem can be approached with portfolio optimization algorithms such as implicit enumeration and dynamic programming. These algorithms have been used for solving the redundancy allocation problem (see references in Kuo and Prasad, 2000; Kuo and Wan, 2007).

In our development, we combine elements of the dynamic programming approach by Liesiö et al. (2007) and the implicit enumeration algorithm by Liesiö (2014) to compute all non-dominated portfolios. This gives a two-phased algorithm in which the first phase implicitly enumerates all possible portfolios and

returns a superset of the non-dominated set, and the second phase removes dominated portfolios by pairwise comparisons. In the computational development, we restrict our analysis to coherent systems.

The reliability function $r(p)$ of coherent systems is monotone in its arguments. Without loss of generality, we assume that $r(p)$ is decreasing in its arguments (if $r(p)$ were to be increasing with respect to p_i, we could consider the complement of event i, which probability would be $q_i = 1 - p_i$, thus making $r(p)$ decreasing with respect to q_i). We also assume that $r(z; p)$ is increasing in z, because if the change from $z_i = 0$ to $z_i = 1$ would lead to lower reliability, then this improvement can be excluded from the analysis, because no decision maker would select such an alternative. Under these assumptions, we develop an implicit enumeration algorithm for solving problem (2).

The basis for the implicit enumeration algorithm is in Table 1. It consists of two distinct phases. The first phase on lines 1-22 is the implicit enumeration phase which yields a superset of the non-dominated set. The second phase on line 23 removes dominated portfolios from the superset given by the first phase such that only non-dominated portfolios remain. The details of the algorithm are as follows.

On line 1, the current solution z, and the active alternative k are initialized. Lines 4-8 and 16-22 of the algorithm enumerates all possible portfolios using a last-in-first out ordering (see, e.g. Liesiö, 2014; Martello and Toth, 1990). On line 9, the feasibility of the portfolio is checked; if this is the case, then on line 10, the portfolio is added to the set of potential non-dominated portfolios Z. On line 11, the lowest reliability that the portfolio can have is computed and compared to the reliability of the so called minmax portfolio, which is the portfolio with the lowest reliability lower bound of the portfolios found so far. If the lowest bound is higher, then the current portfolio is updated to this position.

Lines 13-15 form the bounding phase involving two types of bounds. The first type of bounding is based on feasibility. That is, if at an iteration in the `Forward loop` for some k we have that $\sum_{i=1}^{k} c_i z_i \geq B$, then every further iteration made in the `Forward loop` corresponds to an infeasible solution until a Back-track

step is made such that $z_k \leftarrow 0$. Thus, it not necessary to explicitly account for these infeasible portfolios. Instead, the `Forward-loop` can be terminated prematurely and the `Backtrack-step` can be taken instead.

The second bound is based on strict dominance: Assume that for a given k, portfolio z^* strictly dominates every feasible portfolio which is of the form $(z_1, \ldots, z_k, x_{k+1}, \ldots, x_m)$ where $x_j \in \{0, 1\}$. Then the `Forward loop` will yield only portfolios that are dominated until a `Backtrack-step` $z_k \leftarrow 0$ is made. Thus, it is possible to do this `Backtrack-step` immediately without losing any non-dominated portfolios. Note that z^* is initialized as the empty set, wherefore we define that $r(\emptyset, \underline{p}) = -\infty$ for cases where the starting portfolio $(0, \ldots, 0)$ is infeasible.

Before executing line 23, the first part of the algorithm yields the set Z that is guaranteed to contain all non-dominated portfolios. In the second phase, pairwise comparisons of portfolios in Z are carried out to remove those dominated portfolios that were not screened in the bounding phase. After this second phase, only non-dominated portfolios remain.

The above steps are straightforward to implement except the second bounding criterion on line 14; we discuss this in more detail in the following section.

3.4. Computation of bounds

The second bounding criterion in the algorithm in Table 1 is

$$r(z^*; p) > r(z_1, \ldots, z_k, x_{k+1}, \ldots, x_m; p) \forall p \in \mathcal{P}_F, (z_1, \ldots, z_k, x_{k+1}, \ldots, x_m) \in \mathcal{Z}_F \, .$$

Checking this inequality is a combinatorial problem that can be solved by checking $r(z, p) - r(z', p) \geq 0$ for all $p \in \mathcal{P}_F$ and $z' = (z_1, \ldots, z_k, x_{k+1}, \ldots, x_m) \in \mathcal{Z}_F$ as described in the previous section. However, this makes it necessary to evaluate 2^{m-k-1} portfolios, and therefore an exact evaluation as a part of an already exponential time algorithm is not computationally viable. Instead, we derive a computationally inexpensive condition that is sufficient (but not necessary) for the second condition to hold.

Table 1: Algorithm: NonDominatedSetDepthFirst

Input Monotonic dominance relation $\succ$ over portfolios $z \in \mathcal{Z}_F \subseteq \{0,1\}^m$.

Output Set of all non-dominated portfolios ($\mathcal{Z}_N$).

$z \leftarrow (0,\ldots,0)$, $z^* = \emptyset$, $Z \leftarrow \{\}$, $k \leftarrow 0$
if $\sum_i c_i z_i \leq B$ **then**
    $Z \leftarrow \{z\}$
repeat
    `Forward-loop`
    repeat
        $k \leftarrow k+1$
        $z_k \leftarrow 1$
        if $\sum_i c_i z_i \leq B$ **then**
            $Z \leftarrow Z \cup \{z\}$
            if $r(z;\underline{p}) > r(z^*;\underline{p})$ **then**
                $z^* \leftarrow z$
        `Bounding phase`
        if $\sum_i c_i z_i \geq B$ **or**

$$z^* \succ (z_1,\ldots,z_k,x_{k+1},\ldots,x_m) \forall p \in \mathcal{P}_F, (z_1,\ldots,z_k,x_{k+1},\ldots,x_m) \in \mathcal{Z}_F$$

        then
            break loop
    until $k = m$
    `Backtrack-step`
    $z_m \leftarrow 0$
    $k \leftarrow \max(\{j \in \{1,\ldots,m\} : z_j = 1\} \cup \{0\})$
    if $k > 0$ **then**
        $z_k \leftarrow 0$
until $k = 0$
Remove from Z dominated solutions by pairwise comparison.

Consider the function

$$g(x;p) = g(x_{k+1},\ldots,x_m;p) = r(z^*;p) - r(z_1,\ldots,z_k,x_{k+1},\ldots,x_m;p) \ .$$

This function is multilinear in p and x. It is decreasing in x, because improving the reliability of a component will improve the reliability of a coherent system. Because it is multilinear, the function $g(x;p)$ can be factored in the form

$$g(x;p) = a(x_{-i};p)x_i + b(x_{-i};p) \ ,$$

where

$$a(x_{-i};p) = \sum_{\{I\subseteq\{1,\ldots,m\}|i\in I\}} a_I(p) \prod_{j\in I\setminus\{i\}} x_j$$

and

$$b(x_{-i};p) = \sum_{\{I\subseteq\{1,\ldots,m\}|i\notin I\}} a_I(p) \prod_{j\in I} x_j$$

are multilinear functions and $x_{-i} = (x_{k+1},\ldots,x_{i-1},x_{i+1},\ldots,x_m)$, $i = k+1,\ldots,m$. This form shows that the change from $x_i = 0$ to $x_i = 1$ changes the function value by $a(x_{-i};p)$.

We next derive a lower bound on $a(x_{-i};p)$. Any coherent system can be described through a sequence of subsystems (fault tree) including only AND and OR gates. For these gates, it is straightforward to show that the total increase in system reliability achieved through joint improvements in the reliability of two or more of the components is less than the sum of reliability improvements attained by improvements of these individual components. Hence, by induction, this property holds for any coherent system. This implies that the coefficient $a(x_{-i};p)$ is minimal when $x_{-i} = (0,\ldots,0)$. Thus, it is of interest to establish a lower bound for $a(0,\ldots,0;p) = a_\emptyset^{-i}(p)$. Because $a_\emptyset^{-i}(p)$ is multilinear with respect to p, it follows that $a_\emptyset^{-i}(p) = \sum_{I\subseteq\{1,\ldots,n\}} a_I^{-i} \prod_{j\in I} p_j \geq$ for some unique coefficients a_I^{-i}. A lower bound on $a_\emptyset^{-i}(p)$ can be computed analogously to the bound discussed in the beginning of Section 3 as

$$\underline{a}_\emptyset^{-i} := a_\emptyset^{-i} + a_+^{-i}(\overline{p}) + a_-^{-i}(\underline{p}) \ , \tag{3}$$

where $a_{\emptyset}^{-i}$ is the constant, $a_{+}^{-i}(p)$ is the sum of the terms with positive coefficients a_I^{-i} and $a_{-}^{-i}(p)$ is the sum of the terms with negative coefficients a_I^{-i}.

Typically not all of the x_i can be assigned a value of 1 because of the budget constraint. To obtain a better bound, we use a continuous knapsack problem for bounding the lowest value that the simultaneous change of multiple x_i:s can have. The remaining budget is given by $B - \sum_{i=1}^{k} c_i z_i$. We seek to show that regardless of how the $x_i, i = k+1, \ldots, n$ are chosen, the inequality $g(x;p) \geq 0$ holds. Now, consider the following continuous knapsack problem

$$\begin{aligned} \min_{x} \quad & \sum_{i=k+1}^{m} \underline{a}_0^{-i} x_i \\ \text{subject to} \quad & \sum_{i=k+1}^{m} c_i x_i \leq B - \sum_{i=1}^{k} c_i z_i \\ & 0 \leq x_i \leq 1, \ i = k+1, \ldots, m \ . \end{aligned}$$

This standard optimization problem can be solved efficiently in pseudopolynomial time (see Section 2 in Martello and Toth (1990) for details). The optimal value $\tau \leq 0$ of this problem gives a bound on how much at most the value of $g(x,p)$ can be reduced when changing x_i from 0 to 1 under a budget constraint. Thus $r(z^*;p) - r(z';p) \geq r(z^*;p) - r(z;p) + \tau$. We summarize the implications of this discussion in the following lemma.

Lemma 1. *Consider portfolios z^* and z such that $B - \sum_{i=1}^{k} c_i z_i \geq 0$ and $z_{k+1}, \ldots, z_m = 0$. Let $\tau \leq 0$ denote the optimal value of the problem*

$$\begin{aligned} \min_{x} \quad & \sum_{i=k+1}^{m} \underline{a}_{\emptyset}^{-i} x_i \\ \textit{subject to} \quad & \sum_{i=k+1}^{m} c_i x_i \leq B - \sum_{i=1}^{k} c_i z_i \\ & 0 \leq x_i \leq 1, \ i = k+1, \ldots, m \ , \end{aligned}$$

where $\underline{a}_{\emptyset}^{-i}$ are given by equation (3) . If $r(z^*;p) - r(z;p) + \tau \geq 0 \forall p \in \mathcal{P}_F$, then

$r(z^*;p) > r(z_1, \ldots, z_k, x_{k+1}, \ldots, x_m; p) \forall p \in \mathcal{P}_F$ *and* $(z_1, \ldots, z_k, x_{k+1}, \ldots, x_m) \in \mathcal{Z}_F$.

Proof. See the discussion above. □

Consider the function

$$g(x) = \sum_{I \subseteq \{1,\ldots,m\}} \alpha_I \prod_{j \in I} x_j$$

and the problem

$$\begin{aligned} \max_x \quad & g(x) && (4)\\ \text{subject to} \quad & \sum_i c_i x_i \leq B && (5)\\ & 0 \leq x_i \leq 1. && (6) \end{aligned}$$

We first prove a property of the optimal solution that was shown by Hoeffding (1956) for symmetric multilinear functions. By factoring with respect to x_i, we get $g(x) = a(x_{-i})x_i + b(x_{-i})$, where $a(x_{-i}) = \sum_{I \subseteq \{1,\ldots,m\}: i \in I} \alpha_I \prod_{j \in I \setminus \{i\}} x_j$ and $b(x_{-i}) = \sum_{I \subseteq \{1,\ldots,m\}: i \notin I} \alpha_I \prod_{j \in I} x_j$ are multilinear functions. Using this factoring, the objective function can be written as

$$g(x) = \frac{1}{m} \sum_i g(x) = \frac{1}{m} \left(\sum_i a(x_{-i})x_i - \sum_i b(x_{-i}) \right) .$$

Assume that the optimal solution to the problem above is x^*. Consider the following related problem

$$\begin{aligned} \max_x \quad & \frac{1}{m} \left(\sum_i a(x^*_{-i})x_i - \sum_i b(x^*_{-i}) \right) \\ \text{subject to} \quad & \sum_i c_i x_i \leq B && (7)\\ & 0 \leq x_i \leq 1. \end{aligned}$$

Clearly, $x = x^*$ is a feasible solution of this optimization problem whose value at x^* is the same as that of the original problem (4)-(5); thus, (7) is a relaxation of the original problem (4)-(6). Next, we prove that the solutions are equivalent: The necessary and sufficient Karush-Kuhn-Tucker (KKT) condition for x being a local optimal solution of the former and latter problem is

$$\frac{\partial f}{\partial x_i} + u_0 \frac{\partial}{\partial x_i} \left(\sum_i c_i x_i - B \right) + \sum_i u_i (x_i - 1) + \sum_i u_{m+i}(-x_i) = 0 \ \forall i$$

for some $u = (u_0, \ldots, u_{2m}) \leq 0$. For both problems, $\frac{\partial f}{\partial x_i} = a(x_{-i})$ and thus the KKT conditions are identical, wherefore x^* (which is the global optimal solution of the original problem) is a local optimal solution of the latter problem. Problem (7) is a linear programming problem, which implies that a local optimum is a global optimum as well. Therefore, x^* is its global optimal solution.

For optimizing the latter problem, the multiplication by $1/m$ and substraction of the constant $\sum_i b(x^*_{-i})$ can be disregarded, wherefore the optimal solution to the original problem is the same as the solution of the continuous knapsack problem

$$\begin{aligned} \max_x \quad & \sum_i a(x^*_{-i}) x_i \\ \text{subject to} \quad & \sum_i c_i x_i \leq B \\ & 0 \leq x_i \leq 1. \end{aligned} \tag{8}$$

The continuous knapsack problem can be solved as follows (see, e.g. Section 2 in Martello and Toth, 1990): Without loss of generality, assume that $a(x^*_{-i}) > 0$, $c_i > 0$, $B > 0$, $\sum_i c_i > B$, $c_i \leq B$ and that the variables are sorted such that $a(x^*_{-i})/c_i \geq a(x^*_{-(i+1)})/c_{i+1}$ for all i. If these assumptions do not hold, then the problem can be transformed using simple operations to an equivalent knapsack problem where these assumptions hold, or the solution (or the solution with regard to some variables) is trivial.

The critical item x_j is defined by the index $j = \min\{k : \sum_{i=1}^{k} c_i > B\}$. Then the continuous knapsack problem has an optimal solution given by $x_i = 1$ for $i < j$, $x_i = 0$ for $i > j$ and $x_j = B - \sum_{i<j} c_i x_i$. The solution of the continuous knapsack problem is unique if the set $J = \{i : a(x^*_{-i})/c_i = a(x^*_{-j})/c_j\}$ satisfies $J = \{j\}$, that is, there are no other variables that has the same benefit to cost ratio as the critical item. Otherwise, the solution is not unique and all possible solutions are given by $x_i = 1$ for $i \in \{1, \ldots, j-1\} \backslash J$, $x_i = 0$ for $i \in \{j, \ldots, m\} \backslash J$ and $x_i \geq 0$, $\sum_{i \in J} c_i x_i = B - \sum_{k=1}^{\max(\{1,\ldots,j-1\} \backslash J)} c_k x_k$ for $i \in J$.

The solution of the original problem is unique if problem 8 (and hence also problem (7)) has a unique solution and consists of at most one variable that is

not 0 or 1.

The above result characterizes the optimal solution in way that resembles Theorem 2 (i) and its Corollary 2.1 by Hoeffding (1956) for symmetric multilinear functions. The results of Hoeffding was used by Liesiö (2014) to bound the value of a symmetric multilinear function in an implicit enumeration algorithm, in way similar to our method.

Problem (8) helps establish bounds for the branch-and-bound algorithm. Although we do not know $a(x_i^*)$ before solving the problem, we are able to find easily computable bound on these coefficients. This bounding is based on replacing $a(x_i^*)$ in (8) with $\overline{a}_i \geq a(x_i^*)$, wherefore the problem of this new solution will give an upper bound of the optimal value of the knapsack problem. This optimal value of the knapsack problem will then give a bound on how much the objective function can increase. Thus evaluating $g(x)$ at $x = 0$, gives the most conservative bound.

In some instances, the proposed bound may be weak, as illustrated by the following example. Assume that $p_1, p_2 \in [0.8, 0.95]$ and $g(p_1, p_2) = 1 - (1 - p_1)(1 - p_2) = p_1 + p_2 - p_1 p_2$. Here, computing an upper bound using the last form yields the bound $0.95 + 0.95 - 0.8 \cdot 0.8 = 1.26$ which is greater than the trivial bound one implied by the fact that g is a probability. However, applying the bounds on $g(q_1, q_2) = 1 - q_1 q_2$, where $q_i = (1 - p_i), i = 1, 2$, and applying it yields a bound $1 - (1 - 0.95)^2 = 0.9975$ which is strictly less than one. Thus, the efficiency of the bound depends on how the function is represented. Different forms for multilinear functions can be found in Marichal (2014), for instance. In general, the bound is better, the smaller the values of the nonlinear terms of $g(p)$ are. In particular, this bound can be expected to work well if p_i are close to zero (which is often the case in reliability analysis).

4. Computations and results

This section presents numerical computations which were carried out using a standard laptop with 8 GB RAM, 64-bit Windows 7, and Intel(R) Core(TM) i5-

4300U CPU @ 1.90GHz processor with two cores. We did not use parallelization techniques, and consequently only a single core was used.

In our computational analysis, we consider the same residual heat removal system of a nuclear power plant as Toppila and Salo (2013). In this system, there are 31 failure events whose probabilities are within the lower and upper bounds defined by 90 % confidence intervals. The system is characterized by 147 minimal cut sets which contain at most three events each. The minimal cut sets are visualized in Figure 1 (for details about the minimal cut sets and interval-probabilities, see Toppila and Salo, 2013). The system can be improved by eliminating the event, that is reducing the probability of component failure zero. The portfolio of improvements is constrained by how many events can be eliminated. In the modeling framework presented in Section 3, these are modeled by setting $a_i = 0$, $c_i = 1$, and B is the number improvement that can be taken simultaneously.

We solved the non-dominated portfolios for this problem for portfolios that included 1-5 simultaneous improvements. The core indices based on these non-dominated sets are shown in Figure 1. In this figure, the core indices are coded: Black indicates that the component is a core component, white indicates that the component has core index 0 (exterior component), and gray indicates that the component has a core index between 0 and 1(borderline component). We were not able to solve the problem in reasonable time in cases where the optimal portfolio included 6 or more simultaneous improvements. The explanation can be seen from Table 2: The first part of the algorithm, which implicitly enumerates all portfolios, generated almost 15000 non-dominated portfolios in about an hour. To remove dominated portfolios, second part of the algorithm removes dominated portfolios by pairwise dominance checks, which in this case yields in worst case about $(15000^2 - 15000)/2 = 112492500$ comparisons.

The results can be interpreted as follows. From Figure 1, it can be seen that when only one improvement can be included, then only improvements 1-5 have a non-zero core index and they are all borderline, indicating that any of them could be selected. When the at most two improvements can be selected, then

Table 2: Computations of the residual heat removal system.

Number of improvements in portfolio	1	2	3	4	5	6
Number of feasible portfolios	31	465	4495	31465	169911	736281
Number of potential non-dominated portfolios	5	70	242	2265	6764	14609
Number of minmax potential ND portfolios	5	43	173	1592	5794	-
Number of non-dominated portfolios	5	9	21	58	133	NA

Table 3: Computations of the residual heat removal system with narrower intervals.

Budget	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
#full & feasible portf.	31	465	4.5k	31.5k	169.9k	0.7M	2.6M	7.9M	20.2M	44.6M	84.7M	141M	206M	265M	300M
#pot. ND portfolios	3	3	5	32	169	80	99	203	344	280	122	115	116	259	205
#ND portfolios	3	2	5	7	8	12	24	33	29	23	18	16	18	1	1

also improvements 6 and 11 become borderline. When three to five improvements can be selected, then more improvements become borderline. The most significant change is that improvement 1 becomes core. Thus as a summary, if only a few improvements can be selected, then no other improvements become core, which suggests that improvement 1 should have the highest priority regardless of which other improvements are selected.

When the budget is larger, there are fewer portfolios which are dominated by the portfolio with the highest worst case improvement in reliability. In a sense, to find a small set of non-dominated portfolios that would dominate as many dominated portfolios as possible becomes harder. The maxmin portfolio is such that it has good values over a wide range of parameter values. This suggests that in order to find good candidates for eliminating dominated portfolios, it is required to find non-dominated portfolios that perform well locally. These can be used in conjunction with the implicit enumeration algorithm or the pairwise comparisons to eliminate in as early phase as possible the dominated portfolios.

Next we compare our analysis to the analysis in Toppila and Salo (2013), where recommendations are derived based on Fussell-Vesely (FV) risk importance measure. In that paper, improvements with regard to components 1-5 are

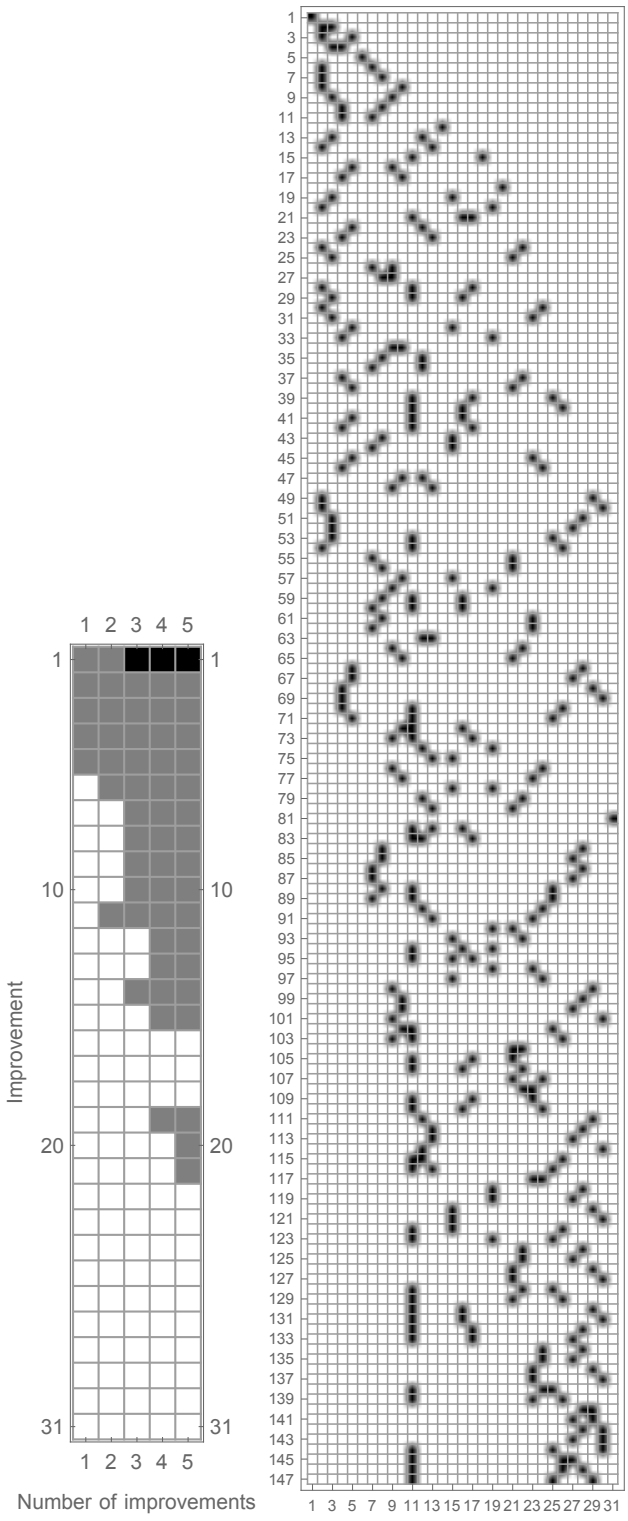

Figure 1: Core indices in the residual heat removal system example (left) and a matrix where each row corresponds to the events that belong to the minimal cut set (right).

Table 4: Computations of the residual heat removal system with short list 10.

Budget	1	2	3	4	5	6	7	8	9	10
#pot. ND portfolios	5	23	35	81	112	94	60	26	8	1
#ND portfolios	5	8	14	27	39	35	21	10	4	1

Table 5: Computations of the residual heat removal system with short list 20.

Budget	1	2	3	4	5	6	7	8
#pot. ND portfolios	5	48	141	859	1995	4229	6862	12965
#ND portfolios	5	9	21	58	132	232	326	NA

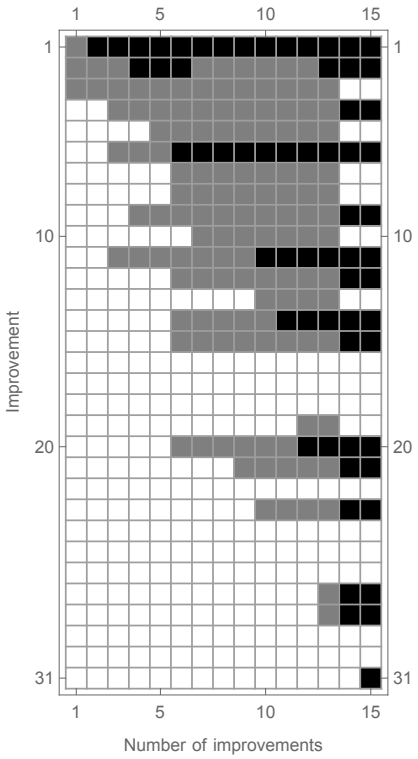

Figure 2: Core indices in the residual heat removal system example with narrower intervals.

identified as non-dominated with respect to all other improvements (to avoid confusion, we abbreviate this as pairwise or pw-ND). When a single improvement can be selected to the portfolio, our method identifies this same set of improvements to be borderline and all others as exterior. However, when more improvements can be added to the portfolio, we see that more improvements become viable: when five improvements can be selected, a total of 15 improvements are borderline or core, compared to the risk measure based analysis, where only improvements 1-5 are pw-ND. We also see that for budgets 3 to 5, improvement 1 is a core improvement. In comparison, the FV measure based analysis does not depend on how many improvements are chosen and hence the decision recommendation does not change either, wherefore the core status of improvement 1 for large budgets is not detected.

We also compute the pw-ND components with respect to the Birnbaum (B) measure using the method by Toppila and Salo (2013). Then the pw-ND improvements are 1, 6, 14, 20, and 31. This ranking is quite different from the ranking based on the FV measure (above), because the B measure does not depend on the probability of the components' failure, wherefore the importance of events that have low probability tends to be exaggerated. In this case, we see that the B measure detects the importance of component 20, which belongs to some ND portfolios at budget level 5. However, as is evident from the figure, also other components than component 20 could be in the ND portfolio.

To solve the ND set with a larger budget, we reduced uncertainties by shrinking the interval endpoints towards the interval mean values so that the interval width was halved. As a result, there are fewer non-dominated portfolios and the range of values that a portfolio can have will become narrower so that the problem becomes easier to solve. With this setting, we could solve the problem to a budget level of 15, corresponding to about 50 % of the events. After this, the reliability function becomes linear with respect to the probabilities. The results are in Figure 2. In this figure, we see a broader perspective on the impact of multiple improvements. At low budget levels, the top three events are borderline or core events (as before in the case with full width intervals,

event 1 becomes core at budget level 2 and above). Event 2 becomes core for budgets 3-6, but loses this status for budgets 7-12. This affect is attributed to the non-linearity of the reliability function, because in a standard knapsack problem, events are taken in the order of highest benefit to cost, wherefore once core, this status should not be lost. An other special case is event 11, which becomes borderline at budget 3 and is the fourth event to become core when the budget is increased.

We also sought to narrow down the number of events that can be included into the portfolio, while keeping the intervals at full width. This effectively reduces the dimension of the search space, but the uncertainties about the probabilities are fully propagated. Thus consider that events 1-10 form a short list from which a portfolio for improvements are selected. All events 1-31 have interval uncertainty about them, but the events that are not on the shortlist cannot be selected for improvement. These results are in Table 4. Here we see that the number of non-dominated portfolios is significantly reduced. Also the number of portfolios found by the implicit enumeration algorithm are very low, which explains why the pairwise comparisons for removing dominated events do not impede solving the problem. We also solved this problem with a short list consisting of events 1-20, but as can be seen from Table 5, the the number of portfolios returned by the implicit enumeration algorithm grew too large for computing the pairwise dominance comparisons when the budget was 8.

We also sought to simplify the problem using the set of portfolios identified by the implicit enumeration algorithm. Clearly, this set contains all non-dominated portfolios and thus, if an event is core or exterior based on core indices computed these sets, then they must be core or exterior in the non-dominated set as well, respectively. Then this problem could be reformulated so that core improvements are included and exterior improvements excluded to/from the portfolio. However, for larger budgets, all improvements were both included and excluded in some portfolios identified by the implicit enumeration algorithm making them borderline, wherefore this technique did not provide any help.

5. Simplifying the reliability function by removing insignificant interactions

Importance measures can be used as a heuristic for solving the redundancy allocation problem (see, e.g. Kuo and Zhu, 2012a,b). Then, by sequentially selecting components of the system for improvement and computing importance measures, one can find good solutions for reliability allocation problems. For instance the Birnbaum importance measure, which is defined as the probability that a basic event i becomes critical to the system, and which can be computed by

$$B_i(p) = \frac{\partial r(p)}{\partial p_i} \ ,$$

can be used for this purpose. From an optimization point of view, this measure linearizes the change with respect to each probability and recommends improving the system in the direct of the steepest accent. This can be modified to take into account the amount of improvement that can be made, the rate of change can be multiplied by the change in that probability, that is $B(p)|p_i^r - p_i| = B(p)\Delta p_i$.

The downside of using the Birnbaum measure for a optimization heuristics is that it does not account for interactions. As stated before, the reliability function is multilinear and, in general, contains products terms that involve interactions. In our method, in which probabilities belong to intervals that can be relatively wide and in which these probabilities are impacted by improvement decisions, the omission of interactions can cause problems.

To overcome this limitation, we use the method of Borgonovo and Smith (2011). In this method, the change of the reliability is considered in *sensitivity cases*: These cases consider simultaneous change in event probabilities such that $p_i \rightarrow p_i + \Delta p_i \ \ \forall i$, and the change is decomposed into three types of interactions measures (Proposition 7, Borgonovo and Smith, 2011) where the measures are

given by the equations

$$\beta_i^1 = B_i(p)\Delta p \tag{9}$$

$$\beta_{i_1,\ldots,i_n}^k = J_{i_1,\ldots,i_n}^k \Delta p_{i_1} \cdots \Delta p_{i_k} \tag{10}$$

$$\beta_i^T = B_i(p)\Delta p_i + \sum_{k=2}^{T} \sum_{i<i_2<\ldots<i_k} J_{i,i_2,\ldots,i_k}^k \cdot \Delta p_{i_1} \cdots \Delta p_{i_k} \ , \tag{11}$$

where

$$J_{i_1,\ldots,i_n}^k = \frac{\partial^k r(p)}{\partial p_{i_1} \cdots \partial p_{i_k}}$$

is the Joint Reliability Importance (see e.g. Kuo and Zhu, 2012a, p. 113). The first components are the linearized changes, as explained in the beginning of this section. The second components are the changes caused by interactions. The third is the Total Order Importance (TOI) measure, as introduced by Borgonovo (2010) (for further development, see Kuo and Zhu, 2012a, pp. 136-137). Note that in this paper, we use the unscaled versions of these measures. The scaled versions can be computed by dividing the measures with the total change related to the sensitivity case.

As noted in Borgonovo and Smith (2011), equations (9)-(11) can be evaluated by computing the reliability function $r(p)$ at $2n + 2$ different points, making the computation of these measure efficient. These measures can be used for decomposing the change of a single probability to its individual impact and to an interaction impact. Here, β_i^1 is a measure of the individual impact of a variable and

$$\beta_i^{\mathcal{I}} := \beta_i^T - \beta_i^1 \tag{12}$$

is a measure of the interaction with other variables.

As discussed in Borgonovo (2010) and Borgonovo and Smith (2011), these measures can be used for reliability functions and models in Probabilistic Safety Assessment that include initiating events. They serve as an opportunity to find the most important interactions that drive the changes in reliability with respect to changes in event probabilities. Notably, if the changes are small, interactions do not play any role at all (Borgonovo and Smith, 2011, Proposition 5).

The TOI help simplify the model by excluding insignificant interaction terms. If the impact of a change in event probability is mainly driven by its individual impact, then all interaction terms in the reliability function regarding this variable can be disregarded. Also, if the total impact of this variable is negligible, then this factor can be possibly fixed as a constant without significantly changing the results. These simplifications make the reliability function more linear and eliminate decision variables, making it easier to find its maximum and minimum values.

5.1. Elimination of insignificant interactions and variables

The elimination procedure proposed by Borgonovo and Smith (2011) is as follows

- Estimate the quantities $\beta_i^T/\Delta f$, $\beta_i^{\mathcal{I}}/\Delta f$ and $\beta_i^1/\Delta f$, where Δf is the total change of the sensitivity case for $i = 1, \ldots, n$
- See if $\beta_i^{\mathcal{I}}/\Delta f << \beta_i^1/\Delta f$. If this holds, then individual effect are the most significant for this variable and they can be neglected. This condition can be checked relying on expert opinion. An automated procedure could use the quantity
$$\epsilon_i = \frac{|\beta_i^{\mathcal{I}}|}{|\beta_i^T|} \, , i = 1, \ldots, n$$
as a decision criterion. Thus if this quantity is below a given threshold, say 10 %, then the interaction could be considered to be negligible, because it would explain only a minor part of the change in that variable.

We build on this framework to simplify the reliability function so that it would be suitable for optimization. To achieve this, we establish guidelines for simplification of the reliability function with respect to those interactions that can be considered negligible.

Consider the function $r(z, p) = r(q)$, where $q_i = p_i(1 - z_i) + ap_iz_i$ as a function of the variable vector $(z, p) = (z_1, \ldots, z_n, p_1, \ldots, p_n)$. For the purposes of dominance, it is of interest to know how the function changes with respect

to the function varying between its lower and upper bounds. Thus for p, it is interesting to compute a sensitivity case where the values of p are changed from $\underline{p}$ to $\overline{p}$.

Next we discuss what sensitivity case should be considered for z. Clearly, changing every z_j variable from 0 to 1 is not that interesting, if a budget constraint would require that only a small subset of the variables can be set to 1. Therefore, we propose that in the sensitivity case, where variables with indices $\ell \in L \subset \{1, \ldots, n\}$ such that $\sum_{\ell \in L} c_i \leq B$ are allowed a change from 0 to 1. Here there the set L can be determined by taking variables in the order of a prior belief about which events should be improved until the budget is exhausted. Such an order can be derived using importance measures, such as Birnabaum, Fussell-Vesely, or Risk Achievement Worth, for instance (for details, see e.g. Kuo and Zhu, 2012a).

The total order interactions β_i^I, the individual impacts β_i^1 and interaction impacts $\beta_i^{\mathcal{I}}$ should be computed based on the sensitivity case presented in the previous paragraph. Then, as proposed in the original method, variables are ordered from the most to the least relevant according to $|\beta_i^T|$. Then, a threshold rule such as the one presented above can be used for determining which variables should be simplified either by linearizing their impact or by substituting them with constants.

5.2. Illustrative example

To illustrate our method, we performed this analysis on the residual heat removal system. The decomposition, where the sensitivity case is such that z_i remain unchanged and p_i are changed from their upper bounds to their lower bounds can be found in Table 6. It is illustrated as a Pareto chart in Figure 3 (lower). We see that most of the change in attributed to the three first probabilities. Of these, the impact of changing the first probability is driven by its individual impact, whereas the impact of changing the second and third probabilities are driven by their interactions with other probabilities. We also see that there are some other probabilities that although have small TOI values,

Table 6: Total order interactions of the RHRS case

Variable index	β_i^T	β_i^1	$\beta_i^{\mathcal{I}}$
1	4.42113×10^{-1}	4.42113×10^{-1}	4.21885×10^{-15}
2	4.61616×10^{-2}	3.57635×10^{-1}	-3.11473×10^{-1}
3	4.11311×10^{-2}	3.12343×10^{-1}	-2.71212×10^{-1}
4	6.14858×10^{-3}	5.65826×10^{-2}	-5.0434×10^{-2}
5	5.21867×10^{-3}	4.82103×10^{-2}	-4.29916×10^{-2}
6	2.12518×10^{-2}	2.12518×10^{-2}	0.
7	3.96072×10^{-3}	3.0077×10^{-2}	-2.61163×10^{-2}
8	3.96072×10^{-3}	3.0077×10^{-2}	-2.61163×10^{-2}
9	2.22785×10^{-3}	1.72601×10^{-2}	-1.50323×10^{-2}
10	1.98507×10^{-3}	1.50743×10^{-2}	-1.30892×10^{-2}
11	1.32923×10^{-3}	3.28293×10^{-2}	-3.15×10^{-2}
12	7.96104×10^{-4}	6.16778×10^{-3}	-5.37167×10^{-3}
13	7.09349×10^{-4}	5.38668×10^{-3}	-4.67733×10^{-3}
14	2.42877×10^{-3}	2.42877×10^{-3}	-4.21321×10^{-15}
15	6.20899×10^{-4}	4.81038×10^{-3}	-4.18948×10^{-3}
16	7.59309×10^{-5}	1.48437×10^{-2}	-1.47678×10^{-2}
17	7.23678×10^{-5}	1.40417×10^{-2}	-1.39693×10^{-2}
18	3.12146×10^{-4}	7.80365×10^{-3}	-7.4915×10^{-3}
19	5.53237×10^{-4}	4.20119×10^{-3}	-3.64796×10^{-3}
20	1.51798×10^{-3}	1.51798×10^{-3}	-4.21321×10^{-15}
21	2.54878×10^{-4}	1.97465×10^{-3}	-1.71977×10^{-3}
22	2.27102×10^{-4}	1.72458×10^{-3}	-1.49748×10^{-3}
23	1.59221×10^{-4}	1.23356×10^{-3}	-1.07433×10^{-3}
24	1.4187×10^{-4}	1.07734×10^{-3}	-9.35467×10^{-4}
25	6.69402×10^{-6}	1.39615×10^{-3}	-1.38946×10^{-3}
26	6.33693×10^{-6}	1.31578×10^{-3}	-1.30944×10^{-3}
27	5.54163×10^{-5}	4.29335×10^{-4}	-3.73919×10^{-4}
28	5.54163×10^{-5}	4.29335×10^{-4}	-3.73919×10^{-4}
29	4.93773×10^{-5}	3.74964×10^{-4}	-3.25586×10^{-4}
30	4.93773×10^{-5}	3.74964×10^{-4}	-3.25586×10^{-4}
31	1.34192×10^{-5}	1.34192×10^{-5}	-4.21325×10^{-15}

have large interaction affects.

For simplicity, we assume that implementing an improvement has the effect that the associated failure event cannot occur, wherefore its probability is zero. Analyzing these improvements gives an upper bound on the achievable reliability in the system. If an improvement is made, the corresponding TOI and individual effects disappear.

For generating the sensitivity case, we approximate the best portfolio to consists of the probabilities of the highest ranked events with respect to the Fussell-Vesely importance measure. Thus in this sensitivity case, $z_i = 1$ for $i = 1, \ldots, B$, where B is the budget, the remaining $z_i = 0$, and the probabilities change from their upper bounds to their lower bounds as previously presented.

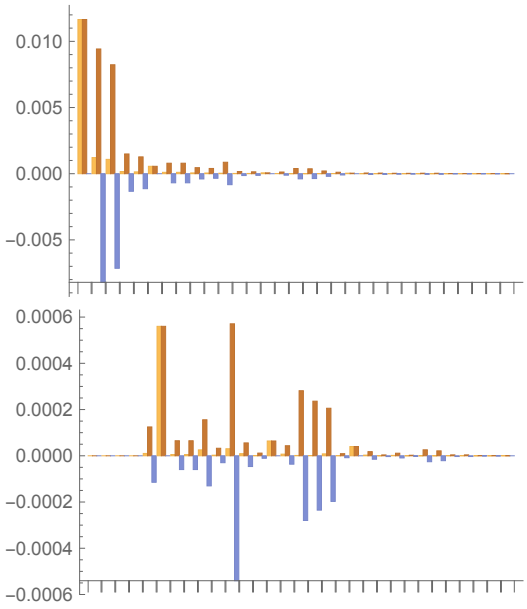

Figure 3: Pareto Charts of TOI (Yellow), Individual impact (Red), and Interaction effect (blue) for two different sensitivity cases. The first is the base case, and the second is the base where events 1-5 have been eliminated.

The TOI decomposition for this sensitivity case is presented in Figure 3 (right). Compared to the initial sensitivity case, we see that the absolute values of the largest TOI values have dropped from about 0.010 to about 0.0006, that is to 6 % of the original value. The improved components have a TOI of zero, because the probabilities related to these events do not change in the sensitivity case. However, we see that interactions start to play a key role much more diversely when the (thought to be) best improvements are done. Thus for more expensive portfolios, these interactions are likely to become significant. The plots of this type was plotted for all budgets up to 10, but they were very similar to this latter sensitivity case.

5.3. Guidelines for simplification

Here, we develop guidelines for simplification that can be done if an interaction is considered insignificant. Consider the following rules:

1. If β_i^T and $\beta_i^{\mathcal{I}}$ are small, then fix this event's probability to the interval upper bound.
2. If β_i^T and β_i^1 are large, and $\beta_i^{\mathcal{I}}$ is small, then fix the value of the probability of all other events to their interval upper bounds in all terms which the event is present.
3. If β_i^T is small but $\beta_i^{\mathcal{I}}$ and β_i^1 are large, then give high priority to this variable in branching.

The rationale for these rules is as follows: If the total impact is small and it does not interact significantly either, then this variable can be effectively treated as a constant. The second rule says that if the interactions are insignificant but the individual impact is significant, then it is better to treat changes on this variable as a linear function, without allowing the probability to interact with other probabilities. The reason for setting the variable to its upper bound stems from the coherence of the system: Approximating the probability of a minimal cut set by setting a probability at its upper bound yields an upper bound for the minimal cut set probability, wherefore the simplified reliability function is a conservative with respect to the original reliability function.

The third rule is motivated by the fact that the bounds proposed to used in the implicit enumeration algorithm and the dominance checking algorithm are based on linearization. Branching on such a variable removes the nonlinearity associated with that variable, which makes the partially evaluated reliability function more linear with respect to the remaining variables. Thus this should lead to better bounds in subsequent branches, hence accelerating the termination of both of these algorithms, and for the implicit enumeration algorithm, reduces the number of dominated portfolios returned by the algorithm. Note that this last rule does not alter the reliability function, only creates a recommendation about how it should be evaluated in the algorithms. Thus if only this rule is used, the analysis is equivalent to the original analysis, and only the computational order is impacted.

An important note about these rules is that it does not matter in which order they are employed. Thus even if they are applied repeatedly and interactively, the end result will be the same. Also, these rules do not have to use the same way for determining if a decomposition term is large or small, which allows for flexibility in tailoring a rule that yields a satisfactory simplification.

These rules should be employed at some budget level, because it is likely that at higher budget levels, the smaller interaction terms are more significant, because the impact potential of the already improved events is reduced (in our RHRS example to zero).

These simplifications reduce the number of interval-valued parameter in the model, wherefore it can also be regarded as a case where more information about the probabilities is elicited, that is intervals are made narrower. This, on the other hand, makes it possible that some new dominances arise if rule 2 is used, which are not present with the unsimplified model. Also, in cases where the variable is fixed in some terms and not in others, the original feasible set is in a sense extended, wherefore it is possible for new non-dominated portfolios to arise. However, these rules try to select for simplification those intervals and terms that according to the TOI measures have least impact, wherefore the impact of the simplification is assumed to be low. Thus, it is hoped that the

non-dominated set of the simplified and unsimplified reliability functions are hopefully close to each other. Note that these extra non-dominated portfolios is easy to remove as a post processing step using the original objective function by pairwise comparison (assuming that this set is reasonably small).

With simplifications, it would be ideal to derive a bound on how much the difference can be. However, the authors are not aware of any reasonable technique which in this case would allow to give a bound on a measure of how similar the sets are. The difficulty stems from the fact that it is not easy to know which portfolios are or even could be non-dominated without effectively generating the set of non-dominated portfolios, which on the other hand would make simplification redundant.

5.4. Computational impact of simplification

We demonstrate the impact of the simplification by two analyses: In the first analysis, we see how much more efficiently the implicit enumeration algorithm can identify the superset that contains the set of non-dominated portfolios, which then again reduces the number of pairwise comparisons needed to remove the dominated portfolios. We compare the non-dominated sets produced by the original and simplified reliability functions and report how large share of these non-dominated sets are equal. Second, we try to solve the instances presented earlier that could not be solved in reasonable time without simplification.

The first case is solved such that it is comparable to the results with narrow intervals, because this instance could be solved for budgets levels up to 50 %. Initially, solving this problem instance for all budget levels took about 30 hours, but after applying the rules that simplify the reliability function, it could be solved in about 2 hours. However, after the simplification, all non-dominated solutions were not found.

A comparison of the simplified method and the exact method can be found in Table 7. In this table, there are the number of potentially ND portfolios and ND portfolios with respect to the simplified reliability function. This ND set corresponds for budget levels 1-5 to the ND set obtained using the original

Table 7: Computations of the residual heat removal system with narrower intervals and simplified reliability function.

Budget	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
#pot. ND portfolios	3	2	5	24	98	27	17	29	39	39	37	25	29	30	35
#ND portfolios (simplified)	3	2	5	7	8	7	6	5	8	6	11	13	6	2	5
Share of actual ND set (%)	100	100	100	100	100	42	75	85	72	74	39	25	67	0	0

reliability function. However, we see that for budgets 6-13, the average share of ND portfolios is from 39%-85 %, whereas for budgets 14-15, the simplified method did not find any ND portfolios.

6. Discussion

This paper has developed a model for solving the redundancy allocation under interval-valued probabilities. In contrast to the reliability optimization models of Sung and Cho (1999), Prasad and Kuo (2000) and Feizollahi and Modarres (2012) who assume more on system structure, our method is, in principle, applicable to any coherent system. However, for our method to be computationally viable, we need to approximate the reliability as the sum of the minimal cut set probabilities. Also, the degrees of redundancy are assumed to be known for each component. In these models, also the degrees of redundancy are solved.

Based on the computational experiments, our method cannot be used for analyzing the full scale fault tree of a nuclear power plant or perhaps not even its larger subsystems. This observation, however, is true even if there is no interval-uncertainty about the probabilities, because the underlying problem to be solved is a knapsack problem with a nonlinear objective function, which belongs to the class of NP-hard problems. Still, our method can provide insights into how uncertainty impacts the prioritization of risk reduction activities. To ensure computational tractability, one could use other methods such as importance measures to narrow the events that can be included to the reduction portfolio

down to a short list of about 10-25 events. Thus, even if some significant events may be missed, at least the events that are on the short list are treated rigorously.

The methods of this paper may be useful in other contexts of optimizing a reliability function with regard to interval-valued probabilities. For instance, consider a case where probabilities are determined by an Weibull distribution, which has interval valued uncertainty about the shape and scale parameters. Then the reliability function is not a multilinear function about the uncertain parameters. However, the method of Borgonovo and Smith (2011) has extensions for analyzing functions that are not multilinear: In these cases, the total order interaction decomposition is, however, only approximate and can be considered to hold only for small changes if the analyzed reliability function is smooth enough.

The width of the intervals was more of a factor in this case than the number of improvement actions that could be selected. This could be seen from the results: Narrowing the intervals made it possible to compute all non-dominated portfolios, while reducing the search space made it possible to solve the problem for only a marginally larger budget.

Importance measures for groups of components is a central topic in reliability research (Kuo and Zhu, 2012b). Possible uses of our method could be to identify sets of components that could be analyzed further with these group importance measures. These groups could be formed from the non-dominated set, because any portfolio can be understood as a binary indicator vector of the events that belong to a group.

In this paper, we consider only a single budget constraint. As discussed in several texts about the implicit enumeration algorithm, inclusion of additional constraints is straightforward. Computationally, this can even be beneficial, because typically additional constraints also reduce the feasible set, wherefore the back-track step due to infeasibility can be done. If the number of constraints is large and they do not reduce the feasible set significantly, then the inclusion of more constraints do not lead to a faster termination of the implicit enumera-

tion algorithm, and consequently checking the validity of these constraints will require additional computational effort.

The developments in this paper have been based on the earlier approach by the authors (Toppila and Salo, 2013), where the reliability of a system is characterized by the minimal cut sets. In the literature, there have been improvements in the representation of the reliability function, which allow more accurate analysis of the reliability of a large system. Especially a representation in terms of a binary decision diagram may be relevant (see e.g. Jung, 2015). The binary decision diagram may serve as a useful data structure for optimization and may allow one to derive tighter bounds for eliminating dominated portfolios, which could improve the computational methods presented in this paper.

When analyzing large problems, there may be so many non-dominated portfolios that enumerating them all would require excessive resources and too long computation time. In these cases, computing a only a subset of all non-dominated portfolios is of interest too, because it still is better to select a non-dominated portfolio than a dominated portfolio. An approximate method could be based on simulation: At simulation round $k = 1, \ldots, N$, fix the probabilities at their lower or upper bound at random (using e.g. a discrete uniform distribution). Compute the solution π_k of problem (2) using these probabilities. Then the core index CI_i of event i can be approximated by computing the fraction

$$CI_i \approx \frac{|\{p \in \{\pi_1, \ldots, \pi_N\} | i \in p\}}{N} \ . \tag{13}$$

A caveat of this method is that it does not find all non-dominated solutions in every problem instance. This is because a non-dominated portfolio may not have the highest reliability for any feasible parameter values. This results in a qualitative loss of the analysis, because the non-dominated portfolios that cannot be found are such that they have a relatively high reliability over a large share of feasible probabilities. Thus, this can be interpreted that these non-dominated portfolios are less affected by epistemic uncertainty about probabilities, wherefore this caveat can be significant in some contexts.

7. Conclusions

We have presented a method for reliability allocation under interval-valued uncertainty about probabilities. The two-phased algorithm consisting of implicit enumeration and pairwise comparison is capable of solving the resulting the set of non-dominated portfolios that characterize reasonable solutions to this problem. However, the computational complexity increases quickly with the size of the problem, wherefore state-of-the-art simplification methods were used. The simplification made it possible to find a significant share of the non-dominated set with considerable less computational effort.

Acknowledgements

Antti Toppila and Ahti Salo acknowledge the financial support of the Nuclear Waste Management Fund (VYR) that was acquired through the Finnish National Nuclear Power Plant Safety Research programme 2011-2014 SAFIR 2014.

Apostolakis, G., 1990. The concept of probability in safety assessment of technological systems. Science 250, 1359–64.

Baraldi, P., Zio, E., Compare, M., 2009. Importance measures in presence of uncertainties. Reliability & Risk Analysis: Theory & Applications 2 (2), 30–41.

Borgonovo, E., 2008. Epistemic uncertainty in the ranking and categorization of probabilistic safety assessment model elements: Issues and findings. Risk Analysis 28 (4), 983–1001.

Borgonovo, E., 2010. The reliability importance of components and prime implicants in coherent and non-coherent systems including total-order interactions. European Journal of Operational Research 205 (3), 485–495.

Borgonovo, E., Smith, C. L., 2011. A study of interactions in the risk assessment of complex engineering systems: An application to space PSA. Operations Research 59 (6), 1461–1476.

Buckley, J., 2003. Fuzzy Probabilities: New Approach and Applications. Vol. 115. Physica-Verlag, Berlin Heidelberg.

Chern, M.-S., 1992. On the computational complexity of reliability redundancy allocation in a series system. Operations Research Letters 11 (5), 309–315.

Coolen, F., 2004. On the use of imprecise probabilities on reliability. Quality and Reliability Engineering International 20 (3), 193–202.

Dempster, A., 1967. Upper and lower probabilities induced by a multivalued mapping. The Annals of Mathematical Statistics 38, 325–339.

Elegbede, A. C., Chu, C., Adjallah, K. H., Yalaoui, F., 2003. Reliability allocation through cost minimization. IEEE Transactions on Reliability 52 (1), 106–111.

Feizollahi, M., Modarres, M., 2012. The robust deviation redundancy allocation problem with interval component reliabilities. IEEE Transactions on Reliability 61 (4), 957–965.

Ha, C., Kuo, W., 2006. Reliability redundancy allocation: An improved realization for nonconvex nonlinear programming problems. European Journal of Operational Research 171 (1), 24–38.

Helton, J., Johnson, J., Oberkampf, W., 2004. An exploration of alternative approaches to the representation of uncertainty in model predictions. Reliability Engineering & System Safety 85 (1-3), 39–71.

Hoeffding, W., 1956. On the distribution of the number of successes in independent trials. The Annals of Mathematical Statistics, 713–721.

Jung, W. S., 2015. A method to improve cutset probability calculation in probabilistic safety assessment of nuclear power plants. Reliability Engineering & System Safety 134 (0), 134 – 142.

Kuo, W., Prasad, V. R., 2000. An annotated overview of system-reliability optimization. IEEE Transactions on Reliability 49 (2), 176–187.

Kuo, W., Wan, R., 2007. Recent advances in optimal reliability allocation. In: Computational Intelligence in Reliability Engineering. Springer, pp. 1–36.

Kuo, W., Zhu, X., 2012a. Importance Measures in Reliability, Risk, and Optimization: Principles and Applications. John Wiley & Sons.

Kuo, W., Zhu, X., 2012b. Some recent advances on importance measures in reliability. IEEE Transactions on Reliability 61 (2), 344–360.

Le Duy, T., Vasseur, D., Couplet, M., Dieulle, L., March 13-17 2011. Parameter and model uncertainty analysis using Dempster-Shafer theory in nuclear probabilistic risk assessment. In: ANS PSA 2011 International Topical Meeting on Probabilistic Safety Assessment and Analysis.

Le Duy, T., Vasseur, D., Dieulle, L., Bérenguer, C., Couplet, M., 2010. Treatment of epistemic uncertainties with non-probabilistic approaches in applications of probabilistic risk assessment. In: Proceedings of PSAM 10 Topical Meeting, Seattle US.

Liesiö, J., 2014. Measurable multiattribute value functions for portfolio decision analysis. Decision Analysis 11 (1), 1–20.

Liesiö, J., Mild, P., Salo, A., 2007. Preference programming for robust portfolio modeling and project selection. European Journal of Operational Research 181 (3), 1488–1505.

Liesiö, J., Mild, P., Salo, A., 2008. Robust portfolio modeling with incomplete cost information and project interdependencies. European Journal of Operational Research 190 (3), 679–695.

Limbourg, P., Savić, R., Petersen, J., Kochs, H., 2008. Modelling uncertainty in fault tree analyses using evidence theory. Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability 222 (3), 291–302.

Marichal, J.-L., 2014. Structure functions and minimal path sets. ArXiv e-prints2014arXiv1401.0803M.

Martello, S., Toth, P., 1990. Knapsack Problems. Wiley New York.

Modarres, M., 2006. Risk Analysis in Engineering: Techniques, Tools, and Trends. CRC Press, New York, pp. 223-232.

Prasad, V. R., Kuo, W., 2000. Reliability optimization of coherent systems. IEEE Transactions on Reliability 49 (3), 323–330.

Rausand, M., Høyland, A., 2004. System Reliability Theory: Models, Statistical Methods, and Applications. Vol. 396. John Wiley & Sons.

Sallak, M., Schon, W., Aguirre, F., et al., 2013. Extended component importance measures considering aleatory and epistemic uncertainties. IEEE Transactions on Reliability 62 (1), 49–65.

Shafer, G., 1976. A Mathematical Theory of Evidence. Vol. 76. Princeton University Press, Princeton.

Sung, C., Cho, Y., 1999. Branch-and-bound redundancy optimization for a series system with multiple-choice constraints. IEEE Transactions on Reliability 48 (2), 108–117.

Tillman, F. A., Hwang, C.-L., Kuo, W., 1977. Optimization techniques for system reliability with redundancy? A review. IEEE Transactions on Reliability 26 (3), 148–155.

Toppila, A., Salo, A., 2013. Computational framework for prioritization of events in fault tree analysis under interval-valued probabilities. IEEE Transactions on Reliability 62 (3), 583–595.

Utkin, L., Coolen, F., 2007a. Imprecise reliability: An introductory overview. In: Levitin, G. (Ed.), Intelligence in Reliability Engineering. Vol. 40 of Studies in Computational Intelligence. Springer Berlin / Heidelberg, pp. 261–306.

Utkin, L., Coolen, F., 2007b. Imprecise reliability: An introductory overview. In: Levitin, G. (Ed.), Intelligence in Reliability Engineering. Vol. 40. Springer Berlin / Heidelberg, pp. 261–306.

Walley, P., 1991. Statistical Reasoning with Imprecise Probabilities. Monographs in Statistics and Applied Probability 42. Chapman and Hall, London.

Weichselberger, K., 2000. The theory of interval-probability as a unifying concept for uncertainty. International Journal of Approximate Reasoning 24 (2-3), 149–70.

Zio, E., 2011. Risk importance measures. In: Pham, H. (Ed.), Safety and Risk Modeling and Its Applications. Springer-Verlag, London, pp. 151–196.