A Variance-Based Estimation of the Resilience Indices in the Preliminary Design Optimisation of Engineering Systems Under Epistemic Uncertainty

Carlos Ortega Absil
PhD candidate, Aerospace Centre of Excellence, Mechanical and Aerospace Engineering
University of Strathclyde, James Weir Building, 75 Montrose Street, Glasgow, UK
E-mail: carlos.ortega@strath.ac.uk

Gianluca Filippi
Postgraduate Researcher, Aerospace Centre of Excellence, Mechanical and Aerospace Engineering
University of Strathclyde, James Weir Building, 75 Montrose Street, Glasgow, UK
E-mail: gianluca.filippi@strath.ac.uk

Annalisa Riccardi
Lecturer, Intelligent Computational Engineering Laboratory, Mechanical and Aerospace Engineering
University of Strathclyde, James Weir Building, 75 Montrose Street, Glasgow, UK
E-mail: annalisa.riccardi@strath.ac.uk

Massimiliano Vasile
Professor, Aerospace Centre of Excellence, Mechanical and Aerospace Engineering
University of Strathclyde, James Weir Building, 75 Montrose Street, Glasgow, UK
E-mail: massimiliano.vasile@strath.ac.uk

Summary
This paper presents novel heuristics for the fast conservative approximation of resilience indices in the preliminary design optimisation of engineering systems under uncertainty. Since the uncertain in the early phases of the design process is mainly of an epistemic nature, Dempster–Shafer theory of evidence is proposed as the reasoning framework. The heuristics proposed in this paper are used to partition the uncertainty space in a collection of subsets that is smaller than the full set of focal elements but still provides a good approximation of Belief and Plausibility. Under suitable assumptions, this methodology renders the approximation of the Belief and Plausibility curves cost-effective for large-scale evidence-based models. Its application to the preliminary-design sizing of a small spacecraft solar array under epistemic uncertainty will be demonstrated.

Keywords: *Optimisation, Uncertainty Quantification, Evidence Theory, Optimisation Under Uncertainty, Preliminary Design, Systems Engineering.*

1 Introduction

Optimisation under Uncertainty has become a fundamental approach to improve design and decision making in systems engineering. Modern techniques in computational intelligence together with current computational power allow designers to model increasingly complex systems and give answers to questions that were out of reach a few years ago regarding the optimal operation of such systems in a variety of scenarios. Such an advancement has led to more exactly quantified design margins and reductions in design budgets without compromising robustness or safety. Nonetheless, the challenge remains as more and more complex problems are tackled better and better algorithms are required. In this sense, there is an increasing interest in improving resilience and performance in complex engineering systems. This can be formulated as what we will call a Resilience Optimisation Problem (ROP).

Resilience Optimisation can be defined as the search for solutions that provide an optimal compromise between system performance and resilience when uncertainty in system design, operational conditions or behaviour is taken into account. Resilience has to be understood as the ability of a system to recover from shocks or to endure severe conditions, where shocks and severe conditions are the manifestation of uncertainty. The recovery, however, is conditional to the property of the system to have some operational capacity even when displaced from nominal conditions. With this in mind, one can see Robust Optimisation and Reliability-based Optimisation as two aspects of Resilience Optimisation.

Robust optimisation aims at minimising the impact of uncertainty on the prediction of the value of the quantity of interest while optimising its value at the same time. Once uncertainty is propagated through the system the variability of the budgets with respect to the uncertain parameters is quantified and one can aim for a good trade-off between predicted budget value and budget variability (expectation-variance methods), or the best budget value under the worst conditions possible (worst case approach).

On the other hand, Reliability-Based Optimisation aims at maximising the ability of a system to retain its operational capacity under uncertainty while maximising performance. The usual practice is to optimise the system budgets with a hard constraint on the probability of failure. A more comprehensive option is to simultaneously optimise the budgets and one or more risk indicators that relate to the probability and/or severity of failure.

Of course a mixed robustness and reliability approach can be envisaged if uncertainty impacts both the system budgets and the system's capability to fulfil its requirements.

Most current methodologies to address these issues focus either on the application of margins and safety factors, or on the measurement of statistical moments over a sample of the space of uncertain parameters. Whereas these practices are widespread and can provide relevant information, they cannot incorporate all forms of uncertainty and do not account for imprecision. The uncertainties the designer needs to cope with in the early phases of the design cycle are often associated to a lack of knowledge, sparsity of background data and imprecise modelling of the system requirements, rather than to the occurrence of aleatory events that are due to nature's inherent randomness. In other words, they lay in the realm of Epistemic Uncertainty.[Helton(1997)] In such a situation, the application of system margins is still possible, but will often lead to overly conservative designs, since the process of devising optimal design margins is hitting the same wall of imprecise definitions and sparse background data.

On the other hand, modelling this uncertainty using standard probabilities can be difficult since it requires additional hypotheses on the probability distribution. A more natural way to tackle these engineering problems[Oberkampf and Helton(2002)] is offered by Imprecise Probability theories, and in particular by Evidence Theory (or Dempster-Shafer Theory[Shafer(1976)Dempster(1967)]), which operates on deductions from the available evidence instead of assuming complete knowledge of the probability distribution.

2 Evidence-Theoretic Design

In Evidence Theory, both input and model uncertainty are defined by means of basic probability assignments (*bpa*) associated to elementary propositions in the space of possible events. Being Θ the set of all possibilities, a *bpa* is a function $m : 2^\Theta \to [0,1]$ verifying

$$m(\emptyset) = 0\,,$$
$$\sum_{A\subset\Theta} m(A) = 1\,.$$

There is a one to one correspondence between any $A \subseteq \Theta$ and the proposition *the true value of* θ *is in* A, where θ is the quantity of interest whose true value needs to be determined. In model-based systems engineering, elementary propositions will often take the form of an uncertain quantity being within a set of intervals, i.e.

$$A = \{u \in [a_l, b_l]\}\,, \quad 1 \le l \le L\,,$$

and their associated $bpa_l = m([a_l, b_l])$. Note *bpa* can be associated to potentially overlapping or disjoint intervals as well as to their union, the latter representing a degree of ignorance. If several uncertain variables are taken into account, one will consider propositions of the kind

$$A = \{\mathbf{u} = (u_1, u_2, \cdots, u_{n_u}) \in \prod_{j=1}^{n_u} [a_{l_j,j}, b_{l_j,j}] = H_\mathbf{l}\}\,,$$

where $\mathbf{l} = (l_1, l_2, \cdots, l_{n_u})$ is the multivariate index associated to hyperrectangular domain $H_\mathbf{l}$. This yields

$$\Theta = \{H_\mathbf{l}\,,\ 1 \le l_j \le L_j, 1 \le j \le n_u\}\,,\ |\Theta| = \prod_{j=1}^{n_u} L_j\,.$$

Assuming independent uncertainties, the *bpa* of every such possibility can be computed as the product of the *bpa* of the elementary propositions regarding each u_j,

$$m(H_\mathbf{l}) = \prod_{j=1}^{n_u} bpa_{l_j,j}\,.$$

After combination of several, possibly conflicting, evidence sources[Dempster(1967)Zhang(1994)Sentz and Ferson(2002)], a map of probability masses is thus assigned to all elements in 2^Θ. The Belief (*Bel*) on and Plausibility (*Pl*) of a given proposition $A \subseteq \Theta$ are defined as

$$Bel(A) = \sum_{B|B\subseteq A} m(B)\,,$$
$$Pl(A) = \sum_{B|B\cap A\neq\emptyset} m(B)\,,$$

i.e. $Bel(A)$ collects the probability masses associated to possibilities satisfying A, whereas $Pl(A)$ collects the masses of possibilities not contradicting A. Hence

$$Pl(A) = 1 - Bel(\bar{A})$$

and Belief and Plausibility can be interpreted as the lower and upper bounds, respectively, imposed by the evidence

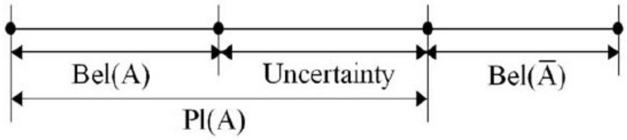

Figure 1: Interpretation of the relation between Belief, Plausibility and (second-order) uncertainty on $P(A)$.

available on the imprecise probability $P(A)$. The difference between $Pl(A)$ and $Bel(A)$ constitutes an indicator of the degree of second-order uncertainty associated to the assessment of $P(A)$. This interpretation is illustrated in Figure 1.

From this point of view, Probability Theory can be viewed as the particular case of Evidence Theory in which Belief and Plausibility converge to the same quantity. Hence, an evidence-theoretic model allows one to treat epistemic and aleatory uncertainty in the same framework.

In the applications that concern this work, the formulation presented translates into considering a mapping of *bpa* over a family of hyperrectangular subsets $H_\mathbf{l}$ of the space of uncertain variables. This family of subsets will from now on, be referred to as U, the uncertainty space, and needs to contain every focal element Ω, this is every subset of Θ with non-null *bpa*:

$$U \supseteq \bigcup \Omega \,,\; \Omega \subset \Theta \,,\; m(\Omega) > 0 \,.$$

The *bpa* structure of U can then be used to calculate the lower (Belief) and upper (Plausibility) bounds on the probability that the value of the quantity of interest $F(\mathbf{u})$ is as expected, e.g. under threshold ν by considering

$$A = \{\mathbf{u} \in U \,|\, F(\mathbf{u}) \leq \nu\} \,,$$

which gives

$$Bel(F(\mathbf{u}) \leq \nu) = \sum_{\overline{\Omega}} m(\Omega) \,,$$
$$Pl(F(\mathbf{u}) \leq \nu) = \sum_{\underline{\Omega}} m(\Omega) \,,$$

with

$$\overline{\Omega} = \{\Omega \subset \Theta | \max_{\mathbf{u} \in H_\mathbf{l} \subseteq \Omega} (F(\mathbf{u})) \leq \nu\} \,,$$
$$\underline{\Omega} = \{\Omega \subset \Theta | \min_{\mathbf{u} \in H_\mathbf{l} \subseteq \Omega} (F(\mathbf{u})) \leq \nu\} \,.$$

Thus, in Robust Design Optimisation, the robustness of a design against the epistemic uncertainty in the system is usually characterised by the curves $Bel(F(\mathbf{u}) \leq \nu)$ and $Pl(F(\mathbf{u}) \leq \nu)$ against ν associated to that design – henceforth referred to simply as *Belief and Plausibility curves*. In particular if F is to be minimised, then A as defined above is the desirable hypothesis, and the robustness index is often chosen as $Bel(F(\mathbf{u}) \leq \nu)$ since it can be interpreted as a conservative estimation of the probability associated to the desirable hypothesis. In Reliability-Based Optimisation, proposition A can instead represent the satisfaction of an operational constraint; the designer will then be interested in the reliability index $Bel(A)$ and its variation against one or more conflicting system budgets.

The drawback of this comprehensive approach for uncertainty quantification is that it leads to an NP-hard problem with a computational complexity that is exponential with the number of epistemic uncertain variables. This is due to the fact that a global maximisation (resp.minimisation) of the quantity of interest is required over each $\Omega \subset \Theta$ having non-null *bpa*.

This work proposes a novel heuristic to produce a progressive approximation of the Belief and Plausibility curves at a reduced computational cost. This approach tries to minimise the estimation error at each iteration and converges faster, under suitable assumptions, to a more precise estimation of the Belief and Plausibility curves than previously proposed partitioning approaches[Vasile et al(2012)Vasile, Minisci, and Wijnands].

Such heuristics can yield huge cut-off in computational resources, allowing one to tackle the complete risk-budget trade-off Pareto front computation for simplified but high-dimensional system models within an affordable time budget. This application will be demonstrated in the last section of this paper by means of the preliminary design optimisation of the solar array of a small spacecraft.

3 Estimation of the Belief and Plausibility Curves

For an exact reconstruction of the Belief (resp. Plausibility) curve, the determination of the worst event (resp. best-case event) is necessary over every subset of the uncertainty space that has a non-null *bpa*. In the general case, this translates into a number of global maximisations (resp. minimisations) of the quantity of interest $F(\mathbf{u})$. This section will focus on the estimation of the Belief curve

$$Bel(F(\mathbf{u}) \leq \nu) = \sum_{\overline{\Omega}} m(\Omega) \,,$$
$$\overline{\Omega} = \{\Omega \subset \Theta | \max_{\mathbf{u} \in H_\mathbf{l} \subseteq \Omega} (F(\mathbf{u})) \leq \nu\}$$

of a design over all possible values of ν. Note the exact computation of the entire curve can be conducted by cumulative sum of *bpa* over the sorted maxima

$$\mathscr{F} = \{\max_{\mathbf{u} \in H_\mathbf{l} \subseteq \Omega} F(\mathbf{u}) \,,\; \Omega \subset \Theta \,,\; m(\Omega) > 0\} \,.$$

The extension to the calculation of Plausibility is immediate, and the ideas exposed hereby can easily be applied to the computation of propositions stated otherwise.

In the algorithm proposed, the whole computation of Belief proceeds by building a tree that has at its root the whole uncertainty space with the associated global worst-case optimisation solution, and at its distal leaves the whole set of focal elements, each one with an associated maximum of the quantity of interest. The heuristic that drives how the tree is built and explored is key to the rapid convergence to the correct Belief and Plausibility values. The overall procedure is schemed in Algorithm 1 and detailed in the following subsections.

3.1 Truncated estimation

The truncated estimation process begins with a global maximisation over the whole uncertainty space U as zeroth iteration

$$S_0 := U\,,\ \bar{F}_0 = \max_{\mathbf{u}\in S_0^0} F(\mathbf{u})\,.$$

This allows to assert

$$\begin{aligned} Bel(F(\mathbf{u}) \le \nu) = 1 &\quad \nu \ge \bar{F}_0\,,\\ Bel(F(\mathbf{u}) \le \nu) \ge 0 &\quad \nu < \bar{F}_0\,, \end{aligned}$$

and is equivalent to propagation of the vacuous Belief function[Dempster(1967)] to quantity F over U.

Then a s-subdivision of the search space is proposed,

$$S_1^1 \cup S_1^2 \cup \cdots \cup S_1^s = S_0^0\,,$$

where s is a hyperparameter of the process. Since this split happens recursively, at iteration $i \ge 1$ one has a set

$$S_i = \{S_i^k\,,\ 1 \le k \le s^i\}$$

of subsets under consideration. Global optimisation is used to obtain the maxima

$$\bar{F}_i = \{\bar{F}_i^k = \max_{\mathbf{u}\in S_i^k} F(\mathbf{u})\,,\ 1 \le k \le m^i\}\,.$$

Let us assume for the sake of simplicity that k is redefined here so that such list is sorted $\bar{F}_i^k \le \bar{F}_i^{k+1}$. Then it stands

$$\mathbf{u} \in H_{\mathbf{l}} \subseteq \bigcup_{\kappa=1}^{k} S_i^\kappa \Longrightarrow F(\mathbf{u}) \le \bar{F}_i^k$$

but not necessarily its reciprocal, which allows to compute the m^i-truncated approximation of the Belief curve

$$\widetilde{Bel}(F(\mathbf{u}) \le \nu) = \sum_{\widetilde{\Omega}} m(\Omega)\,,$$

$$\widetilde{\Omega} = \{\Omega \subseteq \bigcup_{\kappa=1}^{k} S_i^\kappa | \bar{F}_i^k \le \nu\}$$

by cumulative sum of bpa over $\bar{F}_i$. This sum can usually be simplified by considering degenerate bpa structures[Helton et al(2006)Helton, Johnson, Oberkampf, and Sallaberry]. Such approximation is conservative by construction, i.e.

$$Bel(F(\mathbf{u}) \le \nu) \ge \widetilde{Bel}(F(\mathbf{u}) \le \nu)\,,$$

and can indeed be interpreted as a second-order Belief under the evidence provided by $\bar{F}_i$. If the equality holds

$$U = \bigcup \Omega,\ \Omega \subset \Theta,\ m(\Omega) > 0\,,$$

which is usually the case, then

$$Bel(F(\mathbf{u}) \le \nu) = \widetilde{Bel}(F(\mathbf{u}) \le \nu) = 1 \Longleftrightarrow \nu \ge \bar{F}_0\,.$$

Furthermore if it holds that

$$\mathbf{u} \in H_{\mathbf{l}} \subseteq \bigcup_{\kappa=1}^{k} S_i^\kappa \Longleftrightarrow \max_{\mathbf{u}\in H_{\mathbf{l}}}(F(\mathbf{u})) \le \bar{F}_i^k$$

then

$$Bel(F(\mathbf{u}) \le \bar{F}_i^k) = \widetilde{Bel}(F(\mathbf{u}) \le \bar{F}_i^k)\,.$$

Assuming exactitude of the global optimisation, it is clear that $\bar{F}_i \subset \bar{F}_{i+1}$, it will nonetheless be assumed that it is necessary to repeat these optimisations; this assumption will help contain the computational cost and is coherent with the conservative-approximation objective of this work. Thus, the cost of the overall process running for $0 \le i < i_{max}$ is at most $\frac{s^{i_{max}}-1}{s-1}$ global maximisations.

3.2 Heuristics for minimisation of the error

The algorithm proposed in this paper stores in an archive A_S the pairs $(\mathbf{u}, F(\mathbf{u}))$ evaluated by the optimisation process every time it is run over a subdomain S to compute

$$\max_{\mathbf{u}\in S\subseteq U} F(\mathbf{u})\,.$$

The archive can include information of previous iterations too. In the case that a deterministic optimisation algorithm is employed, a pre-sample of U can be used to increase the information available during the first iterations.

After the optimisation, A_S is used to decide on a s-subdivision the current space. A function σ is defined

$$\sigma : A_S \longrightarrow \{S^1, S^2, \cdots, S^s\}\,,\ S^1 \cup S^2 \cup \cdots \cup S^s = S\,.$$

An appropriate choice of σ will lead to the construction of a tree such that it can be truncated at the desired depth with minimum approximation error of the Belief curve.

The heuristics proposed hereby will consider dividing S along one direction of uncertainty u_j at a time. Furthermore, it will be considered that

$$S^k = \bigcup (H_{\mathbf{l}}\,|\, H_{\mathbf{l}} \subset S \wedge l_j = k)\,,$$

which is equivalent to subdivide S along all intervals $[a_l, b_l] \,|\, bpa_{l,j} > 0$ for one of the non-singleton variables $u_j \,|\, L_j > 1$. The statements in 3.1 stand by considering

$$s \ge \max_j(L_j)\,.$$

This will from now on be referred to as *breadth-first* exploration of the truncated estimation tree.

Under such premises, defining σ reduces to selecting the direction u_j along which next split will take place. Since the truncated estimation is conservative by construction, σ is chosen hereby so to partition S along the direction that, according to A_S, captures the highest variability of the system budget with respect to the worst case in S. The idea of systematic partition along the u_j by sensitivity analysis on F is introduced in[Helton et al(2006)Helton, Johnson, Oberkampf, and Sallaberry]. Here we consider, for each non-singleton u_j of S, the list of maxima

$$\widetilde{F} = \{\widetilde{F}^k = \max_{\substack{(\mathbf{u},F(\mathbf{u}))\in A_S \\ \mathbf{u}\in S^k}} F(\mathbf{u})\,,\ 1 \le k \le L_j\}\,,$$

which constitutes a prediction of the next-iteration maxima in S if that direction is selected for subdivision. Let us assume once again that k is redefined so that such list is sorted $\widetilde{F}^k \le \widetilde{F}^{k+1}$. The direction selected will then be

$$u_j \mid j = \arg\max_j \frac{\sum_{k=1}^{L_j-1} (\widetilde{F}^{L_j} - \widetilde{F}^k)^2}{L_j - 1}\,,$$

which, by analogy with a variance measure, gives the *variance-based* designation.

This heuristics is designed as to favour a desirable estimated Belief curve over S, i.e. one that grows slowly in the high robustness values. If the maxima in $\widetilde{F}$ constitute a good approximation of the actual maxima over the S^k, which will be the case if the global optimiser explored S effectively, this will compensate the conservative approximation of the truncated estimation. Otherwise, the possible effects of under-exploration of some regions during the previous global maximisations will be mitigated for subsequent iterations.

It is nonetheless noteworthy that the selection does not account for the *bpa* distribution among the S^k. If subdivisions can be selected that are very heterogeneous in *bpa*, then other σ options are preferable for a fast convergence of $\widetilde{Bel}(F \le \nu)$ to $Bel(F \le \nu)$. The authors propose for instance maximising the area under the next-iteration prediction of the overall curve.

Note also that the purpose is here to obtain a precise approximation of the overall Belief curve at a given cost. The designer might be interested in a higher detail for the pessimistic cases, for example, or be only interested in $Bel(F \le \nu)$ for a given ν; then one should explore the tree otherwise than breadth-first. Combining the ideas exposed in[Vasile et al(2012)Vasile, Minisci, and Wijnands] with a $\sigma(A_S)$ subdivision function to accelerate convergence will be the focus of future research.

4 Preliminary reliability-based design of the solar array of a small spacecraft

4.1 The problem

This section presents the application of the algorithm presented in Section 3 to the reliability-based sizing of the solar array of a small spacecraft power system, to be optimised in terms of construction cost and total power-generating surface. Three different formulations of increasing complexity will be proposed in the following sections, where each one is a particular case of the next.

Algorithm 1 Variance-based breadth-first reconstruction of the truncated Belief curve

Initialise $S = U$, $S_0 = \{S\}$ and $i = 0$
while $i < i_{max}$ **do**
  $\bar{F} \leftarrow \{\emptyset\}$, $S_{i+1} \leftarrow \{\emptyset\}$
  for all $S \in S_i$ **do**
    $\bar{F} \leftarrow \bar{F} \cup \{\max_{\mathbf{u}\in S} F(\mathbf{u})\}$
    $A_S \leftarrow$ global optimization history sample
    $S_{i+1} \leftarrow S_{i+1} \cup \sigma(A_S)$
  end for
  Reconstruct $\widetilde{Bel}$ curve from sorted($\bar{F}$) and *bpa*
  Apply termination condition if any
  $i \leftarrow i + 1$
end while
return Last $\widetilde{Bel}$ curve

In all of them, a design will consist on a certain choice of the quantities:

- $A \in [A_{min}, A_{max}]$, the power-generating surface of the solar panel, [m^2].
- $\mu \in [0,1]$, defines the proportion of cells of type I used in the solar panel. Each type, I and II, has its:
  - Best and worst-case solar efficiencies and failure profile, modeled as expert-provided probability assignments to efficiency intervals.
  - Cost per square meter of power-generating surface.

Hence the construction cost of a design $C(\mu, A)$ can be computed independently of the uncertainties. The sources of risk are, besides the solar cell efficiencies:

- Uncertainty on the power consumption of each of the subsystems, mostly due to lack of definition of the exact mission requirements, modeled as expert-provided probability assignments to power requirement intervals.
- Uncertainty on the power generation, mostly due to sparse background data on components recently adopted by the satellite provider, modeled as expert-provided probability assignments to power efficiency intervals.

The model considers 11 power consumptions of low design margin defined over an only interval, 14 power consumptions of high design margin defined over two intervals with distinct probability assignments, and 6 efficiency power ratios also defined over two intervals. Hence $\dim(U) = 31$ and there are $\delta = 20$ non-singleton directions u_j. The uncertainty space is composed of 2^{20} adjacent focal elements.

Note that it is the epistemic uncertainty that is predominant in all cases at this stage of the design. Hence all uncertainties have been modeled as epistemic and will be propagated through the system model by means of Evidence Theory. The reliability index selected is:

$$Bel(P_{gen} \geq P_{req})$$

Where P_{gen} is the power generated by the solar array and P_{req} is the power level required by the system, both uncertain. Thus, $Bel(P_{gen} \geq P_{req})$ is the most conservative probability estimation associated to the event of *satisfying the power requirements of the system* that can be inferred from the available evidence, and one will be interested in its maximisation or equivalently in the minimisation of the risk index

$$Pl(P_{gen} < P_{req}).$$

4.2 Risk assessment of the worst-case solution

With this formulation the worst-case-scenario optimum d^*_{wcs} is sought for the construction cost $C(\mu,A)$ by solving the problem:

$$\min_{d \in D} C(\mu,A)$$

$$s.t.: P_{gen} \geq P_{req} \quad \forall \mathbf{u} \in U$$

This is equivalent to requesting from the system a reliability index of 100% or risk index of 0% an can be solved analytically in this case by fixing $A|P_{gen} \geq P_{req} \quad \forall u \in U$ and minimising over μ.

For this design solution, the proportion of solar cells of each tipe μ^*_{wcs} is fixed and a risk analysis is then conducted varying the power-generating surface A. The reliability index

$$Bel(P_{gen} \geq P_{req}) = 1 - Pl(P_{gen} < Preq)$$

is presented against $C(\mu^*_{wcs},A)$. This curve is estimated within 11 subsequent iterations of the variance-based algorithm proposed hereby. Since the maximisations are analytical over any subset of focal elements considered, a global optimisation log is not available and is hence mimicked with an initial latin hypersquare sample of cardinality 64. The curves thus obtained are compared to the exact curve computed in an exhaustive fashion requiring maximisation over all the focal elements of uncertainty – 2^{20} analytical maximisations in this case.

By construction this curve acts as a lower bound for the maximum reliability index of any design d^* that lays in the risk-budget Pareto set. In other words, it constitutes a lower bound to the overall reliability-budget trade-off curve whose computation is presented in 4.3. Besides, its rightmost point corresponds to the worst case of the worst-case optimum and is thus assured to belong to the risk-budget trade-off Pareto front.

4.3 Bi-objective formulation

With this formulation the computation of the whole risk-budget trade-off Pareto front is tackled for the construction cost of the solar array $C(\mu,A)$. This can be expressed as:

$$\min_{d \in D} \begin{cases} C(\mu,A) \\ Pl(P_{gen} < P_{req}) \end{cases}$$

This is analogous to solving the family of evidence-based reliability-constrained optimisation problems

$$\min C(\mu,A)$$

$$s.t.: Pl(P_{gen} < P_{req}) \leq \varepsilon$$

$$\varepsilon \in [0,1]$$

The problem above is solved by means of a single run of the multi-objective optimisation algorithm Multi-Agent Collaborative Search (MACS[Zuiani and Vasile(2013)]), using 7 iterations of the variance-based algorithm proposed hereby for the approximation of the risk index at each function evaluation. No additional heuristics are added. Note that this formulation is as of today practically intractable without an approximation method for the risk index even for a problem that allows analytical maximisation over the focal elements, since it would require global optimisation over the design space on top of the exhaustive computation of the index over all the focal elements of uncertainty.

4.4 Three-objective formulation

With this formulation the computation of the whole risk-budget trade-off Pareto front is tackled for the construction cost of the solar array and its power-generating surface simultaneously. This can be expressed as

$$\min_{d \in D} \begin{cases} C(\mu,A) \\ A \\ Pl(P_{gen} < P_{req}) \end{cases}$$

and is analogous to solving the family of bi-objective evidence-based reliability-constrained optimisation problems

$$\min_{d \in D} \begin{cases} C(\mu,A) \\ A \end{cases}$$

$$s.t.: Pl(P_{gen} < P_{req}) \leq \varepsilon$$

$$\varepsilon \in [0,1]$$

The solar array cells are such that type II have lower construction cost per kW of power generated but require higher power-generating surface. This holds both when comparing each type's best-case and worst-case parameters. Hence it is expected to find designs with $\mu = 0$ and $\mu = 1$ at the minimal-budget and minimal-surface extrema of the Pareto front, respectively. This problem is solved with the same set-up described in 4.3.

5 Results

5.1 Risk assessment of the worst-case solution

Figure 2 shows the increasing quality of the estimations obtained in 11 successive iterations of the variance-based approximation algorithm. The convergence to the exact curve on the conservative side is assured by construction of the algorithm, but it is still noteworthy that in this case the convergence rate is large enough as to obtain more precision than is necessary for the purposes that occupy the designer, while achieving a reduction of four orders of magnitude in the computational cost (wrt. its exact computation). The heuristics used constitute a model reduction technique in the sense that they compile information represented along some directions of uncertainty, deemed less relevant. Hence these results are not generic, but the convergence speed will be directly related to the reducibility properties of the index to estimate with respect to the problem uncertain variables in a given probability segment. In other words, the maximum estimation error will be obtained when the effect of every uncertain variable is homogeneous and there is no partitioning more significant than another amongst the considered. The problem defined hereby is found to be dominated by the effect of the uncertainty defined on the 6 power efficiencies, of which only the 3 of them corresponding to cells of type II are relevant hereby. Thus 7 iterations of the algorithm are henceforth deemed sufficient to capture most variability.

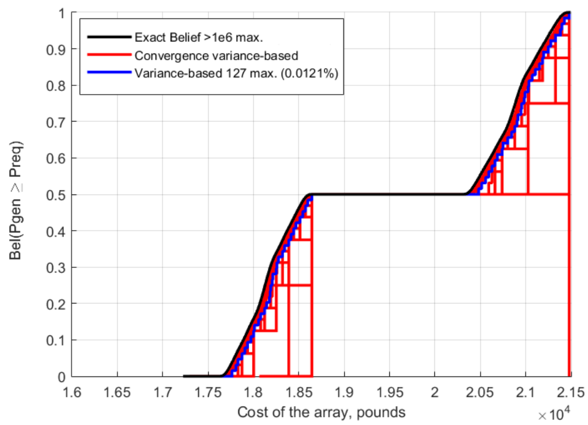

Figure 2: Progressive approximation of the exact reliability-budget curve of a design solution composed entirely of cells of type II ($\mu = 0$). Highlighted, the approximation corresponding to 7 iterations yielding 127 maximisations, i.e. 0.0121% of the computational cost of obtaining the exact curve.

5.2 Bi-objective formulation

Figure 3 illustrates the Reliability Pareto Front obtained for the problem in its bi-objective formulation. As discussed in section 4.2, the curve in figure 2 constitutes a lower bound for the complete reliability-budget trade-off curve and its rightmost point is coincident. In this particular case, since cells of type II have lower construction cost per kW of power generated both in the best and worst case, the leftmost point is also coincident. In this situation one could expect the front to be completely coincident, nevertheless the results show that solving the evidence-based reliability-constraint optimisation problem with a requirement in the reliability index between 0.5 and 0.75 would lead to optimal solutions composed by around 50% of cells of each type.

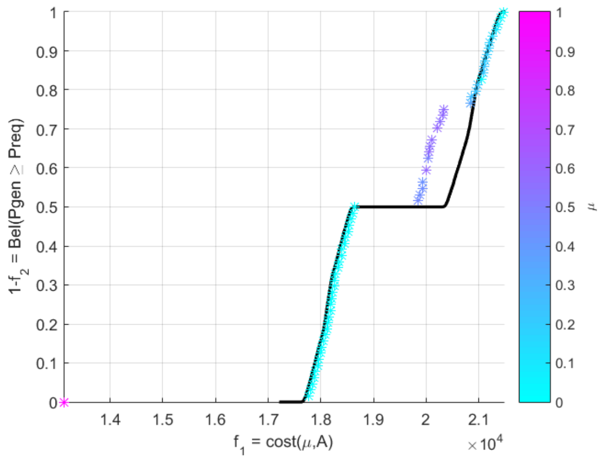

Figure 3: Reliability Pareto front obtained for the design problem in its bi-objective formulation superposed to the exact Belief curve of the worst-case optimum (black line). Colours relate to the proportion of cells of type I and II.

Figure 4 proves that the reliability-budget curve varying A of a solution with $\mu = 0.5$ (dashed line) has both a best case and worst case suboptimal to those of a solution with $\mu = 0$ (solid line), but the former presents two plateaus instead of one and offers thus a higher lower bound on the cdf of the system at a lower construction cost in this reliability range. In this case the designers are more interested in the upper range of reliability and might focus their interest in the budget difference between worst-case cost, nonetheless the availability of this information provides a powerful decision-making tool in a generic scenario.

It can be noted that the quality of the approximation is worse for the design with mixed types of cells using the same estimation set-up. This is due to the fact that each of the maxima used to reconstruct the curve captures, in 7 iterations, the information as divided along 6 of the $\delta = 20$ non-singleton directions of epistemic uncertainty defined. For a design with $\mu = 0.5$, the indices will be more or less equally sensitive to the uncertainty in the parameters of cells of type I and type II, resulting in an homogenisation of the problem landscape. The algorithm, forced to account for more cell-type-related parameters, generates less or no subdivision along the subspace of U corresponding to the uncertainties in the power consumptions, resulting in lower detail. Despite this fact, it has been shown that the quality of the approximation is enough for the bi-objective approach to spot the different behaviour of the solutions and attract attention towards a potentially interesting mixed-type solution.

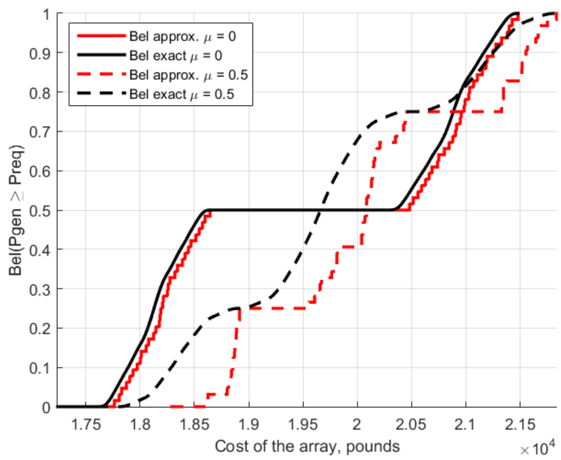

Figure 4: Reliability-budget curves of design solutions with $\mu = 0$ and $\mu = 0.5$. Both the exact curves and those obtained with 7 iterations of the variance-based estimation algorithm are shown.

5.3 Three-objective formulation

Figure 5 shows the family of optimal-budget Pareto fronts obtained in the three-objective formulation for every possible level of reliability requested from the design solution. Of course the fronts with a higher reliability associated are dominated by those that allow a higher risk index. The uppermost front corresponds to the worst-case Pareto optimal solutions. This front can be obtained at a reduced cost using multi-objective worst-case optimisation heuristics such as the ones integrated in MACSminmax[Ortega and Vasile(2017)] . Note that, whereas low and high-reliability solutions constitute almost-linear fronts in the budget space, requesting reliability values between 25 and 75% will lead to more exotically shaped Pareto fronts. In particular, the front becomes non-convex under 50% reliability index, indicating an abrupt change in the properties of the problem landscape. Figure 6 presents the exact same information in a three-dimensional fashion, plus colours relate to the proportion of cells of type I and II used. It can be observed that, as predicted, there is one type of cell that will generally lead to reduction of the cost whereas the other will lead to reduction of the solar array power-generating surface.

6 Conclusions

A methodology has been presented for the fast and conservative estimation of the Belief and Plausibility curves associated to a system budget of quantity of interest. This finds application in Evidence-Theoretic Uncertainty Quantification. The proposed algorithm relies on breadth-first partitioning of the uncertain space and model reduction after analysis of data coming from a global optimisation history. Heuristics to relate the partitioning scheme to the optimisation archive have been proposed and discussed.

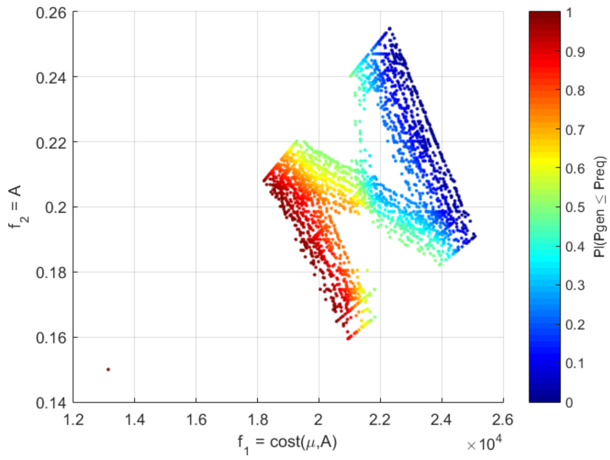

Figure 5: Reliability-budget Pareto Front obtained for the design problem in its three-objective formulation projected to the budget axis, colours relate to the reliability index.

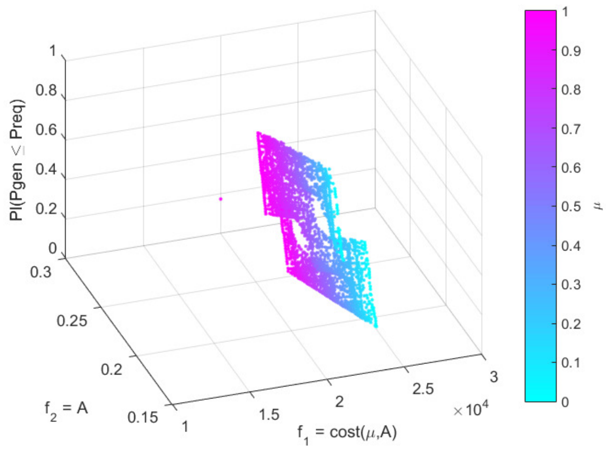

Figure 6: Risk-budgets Pareto front obtained for the design problem in its three-objective formulation, colours relate to the proportion of cells of type I and II.

The overall procedure has been put to the test by means of application in the Expert-Based Reliability Design Optimisation of the solar array of a small spacecraft. For this application, several formulations are proposed of increasing computational complexity to obtain reliability-budget trade-off solutions. Such a detailed analysis is as of today intractable in large-scale engineering problems without a suitable approximation method for the system indices, even if simplified models are in use.

The results show that the proposed methodology can, under suitable model reduction assumptions, provide a large cut-off in computational cost with respect to the exact computation of the Belief and Plausibility curves, while maintaining a minimal approximation error.

It is nonetheless noteworthy that to tackle some of the formulations presented, namely those that involve multi-objective optimisation, only a value of the curve is of interest to drive the search. A preliminary discussion on heuristics to further reduce the cost in such applications has been lead. A broader view on efficient robustness and reliability optimisation algorithms will constitute the focus of future research.

7 Acknowledgement

This work is partially supported by ESTECO Spa. and Surrey Satellite Technologies Ltd. through the project *Robust Design Optimisation of Space Systems* (European Space Agency - Innovation Triangle Initiative).

References

[Dempster(1967)] Dempster AP (1967) Upper and lower probabilities induced by a multivalued mapping. The Annals of Statistics 28:325–339

[Helton(1997)] Helton JC (1997) Uncertainty and sensitivity analysis in the presence of stochastic and subjective uncertainty. Journal of Statistical Computation and Simulation 57:3–76

[Helton et al(2006)Helton, Johnson, Oberkampf, and Sallaberry] Helton JC, Johnson J, Oberkampf WL, Sallaberry C (2006) Sensitivity analysis in conjunction with evidence theory representations of epistemic uncertainty. Reliability Engineering and System Safety 91 (10-11):1414–1434

[Oberkampf and Helton(2002)] Oberkampf W, Helton JC (2002) Investigation of evidence theory for engineering applications

[Ortega and Vasile(2017)] Ortega C, Vasile M (2017) New heuristics for multi-objective worst-case optimization in evidence-based robust design. In: 2017 IEEE Congress on Evolutionary Computation (CEC2017)

[Sentz and Ferson(2002)] Sentz K, Ferson S (2002) Combination of evidence in dempster-shafer theory. In: 6th World Multi-conference on Systemics. Cybernetics and Informatics

[Shafer(1976)] Shafer G (1976) A Mathematical Theory of Evidence. Princeton University Press

[Vasile et al(2012)Vasile, Minisci, and Wijnands] Vasile M, Minisci E, Wijnands Q (2012) Approximated computation of belief functions for robust design optimization. In: 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference

[Zhang(1994)] Zhang L (1994) Representation, independence, and combination of evidence in the dempster-shafer theory

[Zuiani and Vasile(2013)] Zuiani F, Vasile M (2013) Multi-agent collaborative search based on tchebycheff decomposition. Computational Optimization and Applications DOI 10.1007/s10589-013-9552-9