

Buffered environmental contours

K. R. Dahl & A. B. Huseby

*Department of Mathematics,
University of Oslo, Norway*

ABSTRACT: The main idea of this paper is to use the notion of buffered failure probability from probabilistic structural design, first introduced by Rockafellar & Royset (2010), to introduce buffered environmental contours. Classical environmental contours are used in structural design in order to obtain upper bounds on the failure probabilities of a large class of designs. The purpose of buffered failure probabilities is the same. However, in contrast to classical environmental contours, this new concept does not just take into account failure vs. functioning, but also to which extent the system is failing. For example, this is relevant when considering the risk of flooding: We are not just interested in knowing whether a river has flooded. The damages caused by the flooding greatly depends on how much the water has risen above the standard level.

1 INTRODUCTION

Environmental contours are widely used as a basis for e.g., ship design. Such contours allow the designer to verify that a given mechanical structure is safe, i.e., that the failure probability is below a certain value. A realistic model of the environmental loads and the resulting response is crucial for structural reliability analysis of mechanical constructions exposed to environmental forces. See Winterstein et al. (1993) and Haver & Winterstein (2009). For applications of environmental contours in marine structural design, see e.g., Baarholm et al. (2010), Fontaine et al. (2013), Jonathan et al. (2011), Moan (2009) and Ditlevsen (2002).

The traditional approach to environmental contours is based on the well-known *Rosenblatt transformation* introduced in Rosenblatt (1952). This transformation maps the environmental variables into independent standard normal variables. Using the transformed environmental variables a contour with the desired properties can easily be constructed by identifying a sphere centered in the origin and with a suitable radius. More specifically, the sphere can be chosen so that any non-overlapping convex failure region has a probability less than or equal to a desired exceedence probability. The corresponding environmental contour in the original space can then be found by transforming the sphere back into the original space.

Alternatively, an environmental contour can be constructed directly in the original space using Monte Carlo simulation. See Huseby et al. (2013), Huseby et al. (2015a) and Huseby et al. (2015b). Contours constructed using this approach will always be convex

sets. This yields a more straightforward interpretation of the contours. Another advantage of this approach is a more flexible framework for establishing environmental contours, which for example simplifies the inclusion of effects such as future projections of the wave climate related to climatic change. See Vanem & Bitner-Gregersen (2012).

In the present paper we introduce a new concept called buffered environmental contours. This concept is based on the notion of buffered failure probability from probabilistic structural design, first introduced by Rockafellar & Royset (2010). Contrary to classical environmental contours, this new concept does not just take into account failure vs. functioning, but also to which extent the system is failing. For example, this is relevant when considering the risk of flooding: We are not just interested in knowing whether a river has flooded. The damages caused by the flooding greatly depends on how much the water has risen above the standard level.

The structure of this paper is as follows: In Section 2, we recall the classic definition of failure probability in probabilistic structural design and compare this to the concept of buffered failure probability, as defined in Rockafellar & Royset (2010). Furthermore, we recall some of the arguments favoring the buffered failure probability over the regular failure probability. Then, in Section 3, we recall the concept of environmental contours and how such contours are used in structural design in order to find upper bounds on the failure probabilities of a large class of designs. In Section 4, we introduce the new concept of buffered environmental contours, and argue that these contours are better suited than the classical ones in cases where the

level of malfunctioning is important. Finally, in Section 5, we apply the proposed contours to a real life example, and compare the contours to the classical environmental contours.

2 STRUCTURAL DESIGN AND THE BUFFERED FAILURE PROBABILITY

In probabilistic structural design, it is common to define a *performance function*¹ $g(\mathbf{x}, \mathbf{V})$ depending on some design variables $\mathbf{x} = (x_1, x_2, \dots, x_m)'$ and some environmental quantities² $\mathbf{V} = (V_1, V_2, \dots, V_n)' \in \mathcal{V}$, where $\mathcal{V} \subseteq \mathbb{R}^n$. The design variables can be influenced by the designer of the structure, and may represent material type or layout. The quantities are usually random, and cannot be directly impacted by the designer. Hence, they may describe environmental conditions, material quality or loads. To emphasize the randomness of the quantities, we denote them by capital letters. In contrast, the design variables are controlled by the designer and hence denoted by small letters.

For a given design \mathbf{x} , $g(\mathbf{x}, \mathbf{V})$ represents the performance of the structure, and is called the *state of the structure*. A given mechanical structure can withstand environmental stress up to a certain level. The *failure region* of the structure is the set of states of the environmental variables that imply that the structure fails. The performance function is defined such that if $g(\mathbf{x}, \mathbf{V}) > 0$, the structure is *failed*, while if $g(\mathbf{x}, \mathbf{V}) \leq 0$, the structure is *functioning*. Moreover, for a given \mathbf{x} the set $\mathcal{F}(\mathbf{x}) = \{\mathbf{v} \in \mathcal{V} : g(\mathbf{x}, \mathbf{v}) > 0\}$ is called the *failure region* of the structure³.

2.1 The failure probability, reliability and approximation methods

The failure probability, denoted by $p_f(\mathbf{x})$, of the structure is the probability that the structure is failed. That is, $p_f(\mathbf{x}) = P(g(\mathbf{x}, \mathbf{V}) > 0)$. If $f_{\mathbf{V}}(\mathbf{v})$ is the joint probability density function for the random vector \mathbf{V} , the failure probability is given by:

$$p_f(\mathbf{x}) = \int_{\mathcal{F}(\mathbf{x})} f_{\mathbf{V}}(\mathbf{v}) d\mathbf{v}. \quad (1)$$

For a given \mathbf{x} the *reliability*, $R(\mathbf{x})$, of the system is defined as the probability that the system is functioning, i.e.:

$$R(\mathbf{x}) = 1 - p_f(\mathbf{x}) \quad (2)$$

¹The performance function is sometimes called the *limit-state function*.

²Environmental quantities should here be understood in a broad sense. E.g., for marine structures such quantities typically includes wave height and period. For other types of structures, one may consider e.g., material quality, effects of erosion or corrosion as environmental quantities.

³In some papers, such as Huseby et al. (2013), the failed states are defined as the states such that $g(\mathbf{x}, \mathbf{V}) < 0$. This is just a matter of choice of notation.

A classic problem is to compute the reliability of the system. In order to do so, we need to compute the integral (1). In many cases it is difficult to obtain an analytical solution to this. To overcome this issue various approximation methods have been proposed. Two traditional methods for doing this are the *first-order reliability method* (FORM) and the *second-order reliability method* (SORM). The basic idea of the first-order reliability method is to approximate the failure boundary at a specific point by a first order Taylor expansion. The idea behind SORM is similar, but using a second order Taylor expansion instead. In both cases, the approximated failure probability can be used to optimize the structural design, i.e. determine a feasible design which has an acceptable failure probability.

2.2 Return periods

As is common in structural design models, we view \mathbf{V} as representing the average value of the relevant environmental variables in a suitable time interval of length L . Based on this and knowledge of the performance function g it is possible to compute the so-called *return period*. This is done as follows:

We consider the environmental exposure of the given design from time $t \geq 0$. The time axis is divided into intervals of some specified length L , and we let \mathbf{V}_i denote the average environmental quantity in the i th period, $i = 1, 2, \dots$. It is common to assume that $\mathbf{V}_1, \mathbf{V}_2, \dots$ are independent and identically distributed. This is a fairly strict assumption, but as it is so frequently used in structural design, we assume this as well. We then let $T := \min\{i : g(\mathbf{x}, \mathbf{V}_i) > 0\}$. By the assumptions it follows that T is geometrically distributed with probability $p_f = P(g(\mathbf{x}, \mathbf{V}) > 0)$. The *return period* is defined as $E[T] = 1/p_f$. Thus, the return period can be interpreted as a property of the distribution of $g(\mathbf{x}, \mathbf{V})$. Hence, it suffices to analyze this distribution, which is what we will focus on in this paper.

2.3 The buffered failure probability

The approximations made by FORM and SORM can sometimes be too crude and ignore serious risks. Therefore, we will consider the buffered failure probability, introduced by Rockafellar & Royset (2010) as an alternative to the failure probability. This concept relates closely to the conditional value-at-risk (also called expected shortfall, average value-at-risk or expected tail loss), which is a notion frequently used in mathematical finance and financial engineering, see Pflug (2000), Rockafellar (2007) as well as Rockafellar & Uryasev (2000).

Recall that for any level of probability α , the α -quantile of the distribution of a random variable is the value of the inverse of its cumulative distribution function at α . For the random variable $g(\mathbf{x}, \mathbf{V})$, we let

$q_\alpha(\mathbf{x})$ denote its α -quantile. Similarly, for any probability level α , the α -superquantile of $g(\mathbf{x}, \mathbf{V})$, $\bar{q}_\alpha(\mathbf{x})$, is defined as:

$$\bar{q}_\alpha(\mathbf{x}) = E[g(\mathbf{x}, \mathbf{V}) | g(\mathbf{x}, \mathbf{V}) > q_\alpha(\mathbf{x})]. \quad (3)$$

That is, the α -superquantile is the conditional expectation of $g(\mathbf{x}, \mathbf{V})$ when we know that its value is greater than or equal to the α -quantile. Rockafellar & Royset (2010) then define the buffered failure probability, $\bar{p}_f(\mathbf{x})$, as follows:

$$\bar{p}_f(\mathbf{x}) = 1 - \alpha, \quad (4)$$

where α is chosen so that $\bar{q}_\alpha(\mathbf{x}) = 0$. Note that from the previous definitions we have:

$$\bar{p}_f(\mathbf{x}) = P(g(\mathbf{x}, \mathbf{V}) > q_\alpha(\mathbf{x})) = 1 - F(q_\alpha(\mathbf{x})) \quad (5)$$

where F denotes the cumulative distribution function of $g(\mathbf{x}, \mathbf{V})$.




Figure 1: Buffered failure probability calculation where: $p_f(\mathbf{x}) = 0.048$, $q_\alpha(\mathbf{x}) = -0.743$, $\alpha = F(q_\alpha(\mathbf{x})) = 0.879$, and $\bar{p}_f(\mathbf{x}) = 1 - \alpha = 0.121$.

In order to show how to calculate the buffered failure probability $\bar{p}_f(\mathbf{x})$, we consider the plot shown in Figure 1. The curve in the plot represents the cumulative distribution function of the performance function, $g(\mathbf{x}, \mathbf{V})$. As an example we have chosen a Gaussian distribution with mean value -2.5 and standard deviation 1.5 . For this distribution we have $F(0) = 0.952$, as can also be seen in the figure by considering the right-most vertical dashed line starting at 0 on the x -axis, and the corresponding upper horizontal dashed line starting at 0.952 . Hence, we get that $p_f(\mathbf{x}) = 1 - F(0) = 0.048$. In the figure $p_f(\mathbf{x})$ is the distance between 100%-line and the upper horizontal dashed line.

Using e.g., Monte Carlo simulation it is easy to estimate $q_\alpha(\mathbf{x})$, and we find that $q_\alpha(\mathbf{x}) = -0.743$. In the figure $q_\alpha(\mathbf{x})$ is represented by the leftmost vertical dashed line. By following this line until it crosses the cumulative curve, we find that $\alpha = F(q_\alpha(\mathbf{x})) = 0.879$.

Finally, the buffered failure probability is found to be $\bar{p}_f(\mathbf{x}) = 1 - \alpha = 0.121$. In the figure $\bar{p}_f(\mathbf{x})$ is the distance between 100%-line and the lower horizontal dashed line.

It is easy to see that we always have $q_\alpha(\mathbf{x}) \leq 0$, and thus, it follows that $\alpha = F(q_\alpha(\mathbf{x})) \leq F(0)$. This implies that:

$$\bar{p}_f(\mathbf{x}) = 1 - \alpha \geq 1 - F(0) = p_f(\mathbf{x}).$$

Hence, it follows that the buffered failure probability is more conservative than the failure probability. See Rockafellar & Royset (2010) for a detailed discussion of this.

Rockafellar & Royset (2010) present several advantages of using the buffered failure probability instead of the regular failure probability. The following are some of the key arguments:

- In general, the failure probability $p_f(\mathbf{x})$ cannot be computed analytically, and the techniques commonly used to approximate it, such as FORM or Monte Carlo methods, can sometimes ignore serious risks. This makes it problematic to apply standard non-linear optimization algorithms in connection to structure design. In contrast, non-linear optimization algorithms are directly applicable when using the buffered failure probability instead.
- The buffered failure probability contains more information about the tail behaviour of the distribution of $g(\mathbf{x}, \mathbf{V})$ than the failure probability.
- The buffered failure probability can lead to more computational efficiency in design optimization when the performance function $g(\mathbf{x}, \mathbf{V})$ is expensive to evaluate.

The *buffered reliability*, $\bar{R}(\mathbf{x})$, of the structure is defined as $\bar{R}(\mathbf{x}) = 1 - \bar{p}_f(\mathbf{x})$. Since $p_f(\mathbf{x}) \leq \bar{p}_f(\mathbf{x})$, it follows that $R(\mathbf{x}) \geq \bar{R}(\mathbf{x})$. That is, the reliability of the system is greater than or equal to the buffered reliability. Again, this essentially says that the buffered reliability is more conservative than the reliability.

3 ENVIRONMENTAL CONTOURS

Environmental contours are typically used during the early design phases where the exact shape of the failure region is typically *unknown*. At this stage it may not be possible to express a precise functional relationship between a set of design variables \mathbf{x} and the performance of the structure. Instead we skip \mathbf{x} in the notation and let the design options be embedded in the performance function $g(\mathbf{V})$ itself. In particular we denote the failure region simply by \mathcal{F} , while the corresponding failure probability, $P(\mathbf{V} \in \mathcal{F})$, is denoted by $p_f(\mathcal{F})$.

Although \mathcal{F} is unknown, it may still be possible to argue that \mathcal{F} belongs to some known family, \mathcal{E} , of failure regions. As in the previous sections we consider cases where the environmental conditions can be described by a stochastic vector $\mathbf{V} \in \mathbb{R}^n$ with a known distribution. An important part of the probabilistic design process is then to make sure that $P(\mathbf{V} \in \mathcal{F})$ is acceptable for all $\mathcal{F} \in \mathcal{E}$.

In order to avoid failure regions with unacceptable probabilities, it is necessary to put some restrictions on the family \mathcal{E} . This is done by introducing a set $\mathcal{B} \subseteq \mathbb{R}^n$ chosen so that for any relevant failure region \mathcal{F} which do not overlap with \mathcal{B} , the failure probability $P(\mathbf{V} \in \mathcal{F})$ is *small*. The family \mathcal{E} is chosen relative to \mathcal{B} so that $\mathcal{F} \cap \mathcal{B} \subseteq \partial\mathcal{B}$ for all $\mathcal{F} \in \mathcal{E}$, where $\partial\mathcal{B}$ denotes the boundary of \mathcal{B} . This boundary is then referred to as an *environmental contour*. See Figure 2.

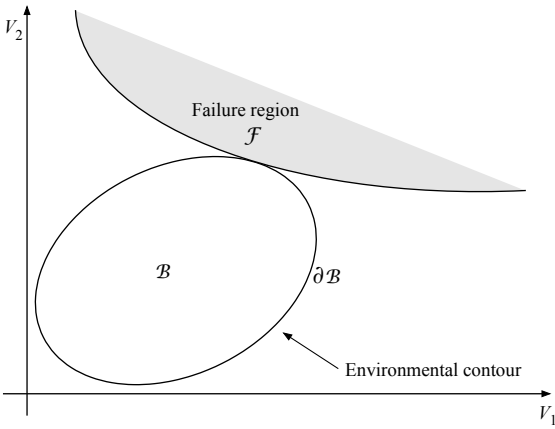


Figure 2: An environmental contour $\partial\mathcal{B}$ and a failure region \mathcal{F} .

Following Huseby et al. (2017) we define the *exceedence probability* of \mathcal{B} with respect to \mathcal{E} as:

$$P_e(\mathcal{B}, \mathcal{E}) := \sup\{p_f(\mathcal{F}) : \mathcal{F} \in \mathcal{E}\}. \quad (6)$$

For a given *target probability* P_e the objective is to choose an environmental contour $\partial\mathcal{B}$ such that:

$$P_e(\mathcal{B}, \mathcal{E}) = P_e$$

We observe that the exceedence probability defined above represents an upper bound on the failure probability of the structure assuming that the true failure region is a member of the family \mathcal{E} . Of particular interest are cases where one can argue that the failure region of a structure is *convex*. That is, cases where \mathcal{E} is the class of all convex sets which do not intersect with the interior of \mathcal{B} . In the remaining part of the paper we will assume that \mathcal{E} satisfies this.

3.1 Monte Carlo contours

There are many possible ways of constructing environmental contours. In this paper we focus on the Monte Carlo based approach first introduced in Huseby et al. (2013), and improved in Huseby et al. (2015a) and Huseby et al. (2015b).

Let \mathcal{U} be the set of all unit vectors in \mathbb{R}^n , and let $\mathbf{u} \in \mathcal{U}$. We then introduce a function $C(\mathbf{u})$ defined for all $\mathbf{u} \in \mathcal{U}$ as:

$$C(\mathbf{u}) := \inf\{C : P(\mathbf{u}'\mathbf{V} > C) \leq P_e\} \quad (7)$$

Thus, $C(\mathbf{u})$ is the $(1 - P_e)$ -quantile of the distribution of $\mathbf{u}'\mathbf{V}$. Given the distribution of \mathbf{V} , the function $C(\mathbf{u})$ can easily be estimated by using Monte Carlo simulation. Thus, let $\mathbf{V}_1, \dots, \mathbf{V}_N$ be a random sample from the distribution of \mathbf{V} . We then choose $\mathbf{u} \in \mathcal{U}$, and let $Y_r(\mathbf{u}) = \mathbf{u}'\mathbf{V}_r$, $r = 1, \dots, N$. These results are sorted in ascending order:

$$Y_{(1)} \leq Y_{(2)} \leq \dots \leq Y_{(N)}$$

Using the sorted numbers we first estimate $C(\mathbf{u})$. Since $C(\mathbf{u})$ is the $(1 - P_e)$ -quantile in the distribution, a natural estimator is:

$$\hat{C}(\mathbf{u}) = Y_{(k)},$$

where k is determined so that:

$$\frac{k}{N} \approx 1 - P_e.$$

Note, however, that this estimator can be improved considerably by using importance sampling. See Huseby et al. (2015b) for details.

For each $\mathbf{u} \in \mathcal{U}$, we also introduce the halfspaces:

$$\Pi^-(\mathbf{u}) = \{\mathbf{v} : \mathbf{u}'\mathbf{v} \leq C(\mathbf{u})\},$$

$$\Pi^+(\mathbf{u}) = \{\mathbf{v} : \mathbf{u}'\mathbf{v} > C(\mathbf{u})\}.$$

We then define the environmental contour as the boundary $\partial\mathcal{B}$ of the *convex set* \mathcal{B} given by:

$$\mathcal{B} := \bigcap_{\mathbf{u} \in \mathcal{U}} \Pi^-(\mathbf{u}) \quad (8)$$

It follows that the exceedence probability of \mathcal{B} with respect to \mathcal{E} is given by:

$$\begin{aligned} P_e(\mathcal{B}, \mathcal{E}) &= \sup\{p_f(\mathcal{F}) : \mathcal{F} \in \mathcal{E}\} \\ &= \sup\{p_f(\Pi^+(\mathbf{u})) : \mathbf{u} \in \mathcal{U}\} \\ &= \sup_{\mathbf{u} \in \mathcal{U}} P(\mathbf{u}'\mathbf{V} > C(\mathbf{u})) = P_e, \end{aligned}$$

where the second equality follows since we have assumed that \mathcal{F} is convex and hence contained in $\Pi^+(\mathbf{u})$ for all $\mathcal{F} \in \mathcal{E}$. In fact for all $\mathbf{u} \in \mathcal{U}$ we have $\Pi^+(\mathbf{u}) \in \mathcal{E}$ as well, and these halfspaces are the maximal sets within \mathcal{E} . Moreover, the last equation follows by the definition of $C(\mathbf{u})$ given in (7). Thus, we conclude that the contour $\partial\mathcal{B}$ indeed has the correct exceedence probability with respect to \mathcal{E} . See Huseby et al. (2017) for further details regarding this.

4 BUFFERED ENVIRONMENTAL CONTOURS

In this section, we introduce a new concept called *buffered environmental contours*. This combines the ideas behind buffered failure probabilities and environmental contours. Before we introduce the main results we review a result on superquantiles which will be essential in our approach (See Rockafellar (2007)).

Proposition 4.1 *Let g_1 and g_2 be two performance functions such that $g_1(V) \leq g_2(V)$ almost surely, and let $\bar{q}_{1,\alpha}$ and $\bar{q}_{2,\alpha}$ denote the α -superquantiles of g_1 and g_2 respectively. Then $\bar{q}_{1,\alpha} \leq \bar{q}_{2,\alpha}$.*

As a corollary of this result we get the following result on buffered failure probabilities:

Corollary 4.2 *Let g_1 and g_2 be two performance functions such that $g_1(V) \leq g_2(V)$ almost surely, and let $\bar{p}_{1,f}$ and $\bar{p}_{2,f}$ denote the buffered failure probabilities of g_1 and g_2 respectively. Then $\bar{p}_{1,f} \leq \bar{p}_{2,f}$.*

For a given performance function g its failure probability, p_f , can be computed based on the failure region of g alone. In contrast, computing the buffered failure probability, \bar{p}_f , requires more detailed information about the distribution of g . We indicate this by expressing \bar{p}_f as a function of g and denoted $\bar{p}_f(g)$.

Just as for classical environmental contours, a *buffered environmental contour* is the boundary $\partial\bar{\mathcal{B}}$ of some suitable set $\bar{\mathcal{B}} \subseteq \mathbb{R}^n$. We shall now describe how the set $\bar{\mathcal{B}}$ can be constructed. As in the previous section we let \mathcal{U} be the set of all unit vectors in \mathbb{R}^n , and let $\mathbf{u} \in \mathcal{U}$. Moreover, we let P_e be a given target probability, and let $C(\mathbf{u})$ be defined by (7). In order to introduce buffering, we let:

$$\bar{C}(\mathbf{u}) := E[\mathbf{u}'\mathbf{V} | \mathbf{u}'\mathbf{V} > C(\mathbf{u})]. \quad (9)$$

Given the distribution of \mathbf{V} , the function $\bar{C}(\mathbf{u})$ can easily be estimated by using Monte Carlo simulation. As in Subsection 3.1, we let $\mathbf{V}_1, \dots, \mathbf{V}_N$ be a random sample from the distribution of \mathbf{V} , and choose $\mathbf{u} \in \mathcal{U}$. Based on the sorted values $Y_{(1)} \leq Y_{(2)} \leq \dots \leq Y_{(N)}$ we first estimate $C(\mathbf{u})$ by $Y_{(k)}$ as previously explained. We then estimate $\bar{C}(\mathbf{u})$ by computing the average value of the sampled values which are greater than $Y_{(k)}$. Thus, we estimate $\bar{C}(\mathbf{u})$ by:

$$\hat{\bar{C}}(\mathbf{u}) = \frac{1}{N-k} \sum_{r>k} Y_{(r)}.$$

For each $\mathbf{u} \in \mathcal{U}$, we also introduce the halfspaces:

$$\bar{\Pi}^-(\mathbf{u}) = \{\mathbf{v} : \mathbf{u}'\mathbf{v} \leq \bar{C}(\mathbf{u})\},$$

$$\bar{\Pi}^+(\mathbf{u}) = \{\mathbf{v} : \mathbf{u}'\mathbf{v} > \bar{C}(\mathbf{u})\},$$

similar to what we did in the previous section. Finally, we define the buffered environmental contour as the boundary $\partial\bar{\mathcal{B}}$ of the *convex set* $\bar{\mathcal{B}}$ given by:

$$\bar{\mathcal{B}} := \bigcap_{\mathbf{u} \in \mathcal{U}} \bar{\Pi}^-(\mathbf{u}) \quad (10)$$

We observe that by (9) we obviously have that $\bar{C}(\mathbf{u}) > C(\mathbf{u})$. By comparing (8) and (10), it is easy to see that this implies that:

$$\mathcal{B} \subset \bar{\mathcal{B}}.$$

Thus, given that the same target probability P_e is used to construct both contours, the buffered environmental contour is more conservative than the classical environmental contour.

The next step is to identify a family \mathcal{G} of performance functions defined relative to the set \mathcal{B} such that $\bar{p}_f(g) \leq P_e$ for all $g \in \mathcal{G}$. We recall that for the classical environmental contour we chose to let \mathcal{E} be the family of all convex failure regions which do not intersect with the interior of \mathcal{B} . Thus, one might think that the natural counterpart for buffered environmental contours would be to let \mathcal{G} be the family of performance functions with convex failure regions which do not intersect with the interior of $\bar{\mathcal{B}}$. In this case, however, we need more control over the distributions of the performance functions. In order to do so we choose $\mathbf{u} \in \mathcal{U}$ and introduce the performance function $\Gamma(\mathbf{u}, \cdot)$ given by:

$$\Gamma(\mathbf{u}, \mathbf{V}) = \mathbf{u}'\mathbf{V} - \bar{C}(\mathbf{u})$$

By (9) we have:

$$\begin{aligned} E[\Gamma(\mathbf{u}, \mathbf{V}) | \Gamma(\mathbf{u}, \mathbf{V}) > C(\mathbf{u}) - \bar{C}(\mathbf{u})] \\ = E[\mathbf{u}'\mathbf{V} | \mathbf{u}'\mathbf{V} > C(\mathbf{u})] - \bar{C}(\mathbf{u}) = 0. \end{aligned}$$

Moreover, by (7) we have:

$$\begin{aligned} \bar{p}_f(\Gamma(\mathbf{u}, \cdot)) &= P(\Gamma(\mathbf{u}, \mathbf{V}) > C(\mathbf{u}) - \bar{C}(\mathbf{u})) \\ &= P(\mathbf{u}'\mathbf{V} > C(\mathbf{u})) = P_e \end{aligned}$$

Since the unit vector \mathbf{u} was arbitrarily chosen, we conclude that the performance function $\Gamma(\mathbf{u}, \cdot)$ has the desired buffered failure probability P_e for all $\mathbf{u} \in \mathcal{U}$.

We will use these performance functions as a basis for constructing the family \mathcal{G} where the $\Gamma(\mathbf{u}, \cdot)$ -functions serve as *maximal* elements in this family. Note that the $\Gamma(\mathbf{u}, \cdot)$ -functions now play a similar role as the halfspaces $\Pi^+(\mathbf{u})$ played in the construction of the family \mathcal{F} . Thus, we let \mathcal{G} be the family of all performance functions g for which there exists a $\mathbf{u} \in \mathcal{U}$ such that $g(\mathbf{v}) \leq \Gamma(\mathbf{u}, \mathbf{v})$ for all $\mathbf{v} \in \mathcal{V}$. By the above discussion the following result is immediate:

Theorem 4.3 For all $g \in \mathcal{G}$ we have $\bar{p}_f(g) \leq P_e$.

Proof: Assume that $g \in \mathcal{G}$. Then there exists a $\mathbf{u} \in \mathcal{U}$ such that $g(\mathbf{V}) \leq \Gamma(\mathbf{u}, \mathbf{V})$ almost surely. Hence, by Corollary 4.2 and the above calculations we have:

$$\bar{p}_f(g) \leq \bar{p}_f(\Gamma(\mathbf{u}, \cdot)) = P_e.$$

□

Having constructed both the set $\bar{\mathcal{B}}$ and the family \mathcal{G} we are now ready to introduce the *buffered exceedence probability* of $\bar{\mathcal{B}}$ with respect to \mathcal{G} defined as:

$$\bar{P}_e(\bar{\mathcal{B}}, \mathcal{G}) := \sup\{\bar{p}_f(g) : g \in \mathcal{G}\}. \quad (11)$$

We note that by the definition of \mathcal{G} it follows that $\Gamma(\mathbf{u}, \cdot) \in \mathcal{G}$ for all $\mathbf{u} \in \mathcal{U}$. Hence, we get:

$$\begin{aligned} \bar{P}_e(\bar{\mathcal{B}}, \mathcal{G}) &= \sup\{\bar{p}_f(g) : g \in \mathcal{G}\} \\ &= \sup\{\bar{p}_f(\Gamma(\mathbf{u}, \cdot)) : \mathbf{u} \in \mathcal{U}\} = P_e, \end{aligned}$$

Thus, we conclude that the contour $\partial\bar{\mathcal{B}}$ indeed has the correct buffered exceedence probability with respect to \mathcal{G} .

If $g \in \mathcal{G}$ and $g(\mathbf{v}) \leq \Gamma(\mathbf{u}, \mathbf{v})$ for all $\mathbf{v} \in \mathcal{V}$, we have:

$$\begin{aligned} \mathcal{F}(g) &\subseteq \mathcal{F}(\Gamma(\mathbf{u}, \cdot)) \\ &= \{\mathbf{v} : \mathbf{u}'\mathbf{v} - \bar{C}(\mathbf{u}) > 0\} \\ &= \{\mathbf{v} : \mathbf{u}'\mathbf{v} > \bar{C}(\mathbf{u})\} = \bar{\Pi}^+(\mathbf{u}) \end{aligned}$$

Thus, the failure region of a performance function $g \in \mathcal{G}$ does not overlap with the interior of the set $\bar{\mathcal{B}}$, but is contained within a halfspace supporting $\bar{\mathcal{B}}$. This is similar to the relation between failure regions in the family \mathcal{E} and the set \mathcal{B} for the classical environmental contours. However, as already pointed out, knowledge about the failure region of a performance function is not sufficient to ensure that the performance function has the correct buffered failure probability.

It may be argued that the choice of the $\Gamma(\mathbf{u}, \cdot)$ -functions as maximal elements in the family \mathcal{G} is too restrictive. In order to have a more flexible framework, it is possible to consider a slightly more general approach where we define:

$$\bar{C}_a(\mathbf{u}) := E[a\mathbf{u}'\mathbf{V} | \mathbf{u}'\mathbf{V} > C(\mathbf{u})] = a\bar{C}(\mathbf{u}), \quad (12)$$

where a is a positive constant. By increasing the a -factor, the contour may be inflated so that it can be used for steeper performance factors.

On the other hand it should be noted that to ensure that a given performance function g has the correct buffered failure probability, it is not necessary that $g(\mathbf{v})$ is dominated by some $\Gamma(\mathbf{u}, \cdot)$ -function for *all* $\mathbf{v} \in \mathcal{V}$. It is sufficient that this holds for \mathbf{v} -values corresponding to the upper tail area of g .

5 NUMERICAL EXAMPLE

In this subsection we illustrate the proposed method by considering a numerical example introduced in Vanem & Bitner-Gregersen (2015). More specifically, we consider joint long-term models for *significant wave height*, denoted by H , and *wave period* denoted by T . A marginal distribution is fitted to the data for significant wave height and a conditional model, conditioned on the value of significant wave height, is subsequently fitted to the wave period. The joint model is the product of these distribution functions:

$$f_{T,H}(t, h) = f_H(h)f_{T|H}(t|h)$$

Simultaneous distributions have been fitted to data assuming a three-parameter Weibull distribution for the significant wave height, H , and a lognormal conditional distribution for the wave period, T . The three-parameter Weibull distribution is parameterized by a location parameter, γ , a scale parameter α , and a shape parameter β as follows:

$$f_H(h) = \frac{\beta}{\alpha} \left(\frac{h - \gamma}{\alpha} \right)^{\beta-1} e^{-[(h-\gamma)/\alpha]^\beta}, \quad h \geq \gamma.$$

The lognormal distribution has two parameters, the log-mean μ and the log-standard deviation σ and is expressed as:

$$f_{T|H}(t|h) = \frac{1}{t\sqrt{2\pi}} e^{-[(\ln(t) - \mu)^2 / (2\sigma^2)]}, \quad t \geq 0,$$

where the dependence between H and T is modelled by letting the parameters μ and σ be expressed in terms of H as follows:

$$\mu = E[\ln(T)|H = h] = a_1 + a_2 h^{a_3},$$

$$\sigma = SD[\ln(T)|H = h] = b_1 + b_2 e^{b_3 h}.$$

The parameters $a_1, a_2, a_3, b_1, b_2, b_3$ are estimated using available data from the relevant geographical location. In the example considered here the parameters are fitted based on a data set from North West Australia. We consider data for two different cases: *swell* and *wind sea*. The parameters for the three-parameter Weibull distribution are listed in Table 1, while the parameters for the conditional log-normal distribution are listed in Table 2. In all the examples we use a return period of 25 years. The models are fitted using sea states representing periods of 1 hour. Thus, we get 24 data points per 24 hours. Thus, the desired exceedence probability is given by:

$$P_e = \frac{1}{25 \cdot 365.25 \cdot 24} = 4.5631 \cdot 10^{-6}.$$

For more details about these examples we refer to (Vanem & Bitner-Gregersen 2015).

Table 1: Fitted parameter for the three-parameter Weibull distribution for significant wave heights

	α	β	γ
Swell	0.450	1.580	0.132
Wind sea	0.605	0.867	0.322

Table 2: Fitted parameter for the conditional log-normal distribution for wave periods

		$i = 1$	$i = 2$	$i = 3$
Swell	a_i	0.010	2.543	0.032
	b_i	0.137	0.000	0.000
Wind sea	a_i	0.000	1.798	0.134
	b_i	0.042	0.224	-0.500

The classical environmental contours are estimated based on the methods presented in Huseby et al. (2013). More specifically, we have used Method 2 presented in this paper. The buffered environmental contours are estimated in exactly the same way, except that $\hat{C}(\mathbf{u})$ is replaced by $\hat{\tilde{C}}(\mathbf{u})$ for all $\mathbf{u} \in \mathcal{U}$.

In Figure 3 and Figure 4 the resulting environment contours are shown. As one expected, the classical environmental contours are located inside their respective buffered contours. Thus, since the target probability P_e is the same for both types of contours, the buffered contours are more conservative than the classical contours.

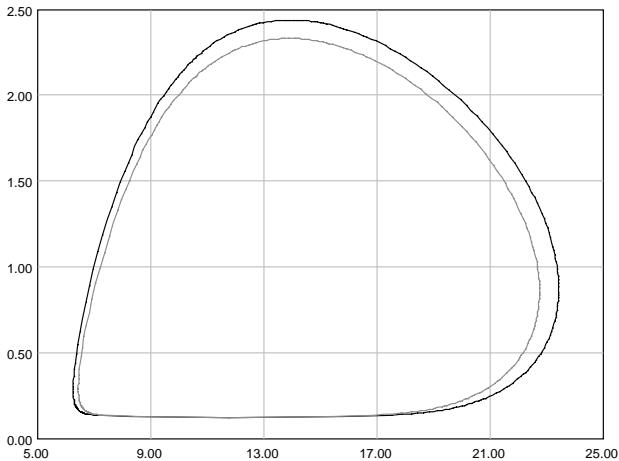


Figure 3: Buffered environmental contour (black) and classical environmental contour (gray) for North West Australia Swell with return period 25 years.

6 CONCLUSIONS AND FUTURE WORK

In the present paper we have introduced the concept of buffered environmental contours, and shown how such contours can be estimated using Monte Carlo simulations. Such contours do not just take into account the probability of failure, but also the consequences of a failure. This is relevant e.g., when




Figure 4: Buffered environmental contour (black) and classical environmental contour (gray) for North West Australia Wind sea with return period 25 years.

analysing the risk of flooding at a given location. While it may not be possible to prevent floodings from occurring, the damage caused by such an event can vary a lot depending on how much the water has risen above the normal level. In some cases only minor damages may be the result. In other cases the consequences can be catastrophic.

For a given target probability, P_e buffered environmental contours are generally more conservative than the classical environmental contours. However, in cases where the consequences are more important than the triggering event itself, a higher target probability might be acceptable as long as the damages are manageable. Thus, in real-life applications a buffered environmental contour may not be so conservative after all. At the same time these contours provide much more information about the tail area of the environmental variables. This may be very useful when a design is optimized.

The buffered environmental contours proposed in this paper are the natural extension of the Monte Carlo contours introduced in Huseby et al. (2013). In particular both contour types are boundaries of convex sets. Sometimes this restriction may lead to contours which include areas of very low probability. Thus, it would be of interest to investigate other ways of constructing buffered contours. In particular, it is possible to modify contours obtained by using the Rosenblatt transformation so that they include buffering. To make this work, however, evaluating the resulting contours becomes very important. The evaluation framework described in Huseby et al. (2017) may serve as a starting point.

Future work in this area also includes the use of buffered environmental contours in design optimization, but with additional design constraints. The question is how such additional constraints can be dealt with. An initial idea is to apply a Lagrange duality method in order to transform the problem into a pre-

viously known form.

It would also be interesting to compare buffered environmental contours to the conservative environmental contours defined by Leira (2008). The contours defined in Leira (2008) are typically larger sets than the environmental contours considered in Section 3, which means that they are more conservative when it comes to classifying structures as safe.

Another idea which requires further investigation is how time can be introduced into this model in a less restrictive way. As mentioned in Subsection 2.2, we consider average stochastic environmental conditions V_1, V_2, \dots over some specified time intervals and assume independence and identical distributions of the V_i 's. A more realistic approach would be to introduce a stochastic process in continuous time modelling the environmental situation. It is interesting to see how this affects the model and what consequences this has for the design optimization.

ACKNOWLEDGEMENTS

This paper has been written with support from the Research Council of Norway (RCN) through the project *ECSADES* Environmental Contours for Safe Design of Ships and other marine structures.

REFERENCES

- Baarholm, G., S. Haver, & O. Økland (2010). Combining contours of significant wave height and peak period with platform response distributions for predicting design response. *Marine Structures* 23, 147–163.
- Ditlevsen, O. (2002). Stochastic model for joint wave and wind loads on offshore structures. *Structural Safety* 24, 139–163.
- Fontaine, E., P. Orsero, A. Ledoux, R. Nerzic, M. Prevesto, & V. Quiniou (2013). Reliability analysis and response based design of a moored fpso in west africa. *Structural Safety* 41, 82–96.
- Haver, S. & S. Winterstein (2009). Environmental contour lines: A method for estimating long term extremes by a short term analysis. *Transactions of the Society of Naval Architects and Marine Engineers* 116, 116–127.
- Huseby, A. B., E. Vanem, & K. Eskeland (2017). Evaluating properties of environmental contours. In M. Cepin and R. Bris (Eds.), *Safety and Reliability, Theory and Applications. Proceedings of the European safety and reliability Conference*, pp. 2101–2109. CRC Press.
- Huseby, A. B., E. Vanem, & B. Natvig (2013). A new approach to environmental contours for ocean engineering applications based on direct monte carlo simulations. *Ocean Engineering* 60, 124–135.
- Huseby, A. B., E. Vanem, & B. Natvig (2015a). Alternative environmental contours for structural reliability analysis. *Structural Safety* 54, 32–45.
- Huseby, A. B., E. Vanem, & B. Natvig (2015b). A new monte carlo method for environmental contour estimation. In T. Nowakowski, M. Mlynczak, A. Jodejko-Pietruczuk, and S. Werbinska-Wojciechowska (Eds.), *Safety and Reliability : Methodology and Applications. Proceedings of the European safety and reliability Conference*, pp. 2091–2098. Taylor & Francis.
- Jonathan, P., K. Ewans, & J. Flynn (2011). On the estimation of ocean engineering design contours. *Journal of Offshore Mech. Arct. Eng* 136(4), 8 pages.
- Leira, B. J. (2008). A comparison of stochastic process models for definition of design contours. *Structural Safety* 30, 493–505.
- Moan, T. (2009). Development of accidental collapse limit state criteria for offshore structures. *Structural Safety* 31, 124–135.
- Pflug, G. (2000). Some remarks on the value-at-risk and the conditional value-at-risk. In S. P. Uryasev (Ed.), *Probabilistic Constrained Optimization. Methodology and Applications.*, Norwell, pp. 272–281. Kluwer Academic Publishers.
- Rockafellar, R. T. (2007). *Coherent Approaches to Risk in Optimization Under Uncertainty.*, Chapter 3, pp. 38–61. INFORMS Tutorials in Operations Research.
- Rockafellar, R. T. & J. O. Royset (2010). On buffered failure probability in design and optimization of structures. *Reliability Engineering and System Safety* 95, 499–510.
- Rockafellar, R. T. & S. P. Uryasev (2000). Optimization of conditional value-at-risk. *Journal of Risk* 2, 21–42.
- Rosenblatt, M. (1952). Remarks on a multivariate transformation. *Ann. Math. Stat.* 23, 470–472.
- Vanem, E. & E. Bitner-Gregersen (2012). Stochastic modelling of long-term trends in wave climate and its potential impact on ship structural loads. *Applied Ocean Research* 37, 235–248.
- Vanem, E. & E. Bitner-Gregersen (2015). Alternative environmental contours for marine structural design – a comparison study. *Journal of Offshore Mechanics and Arctic Engineering* 137, 051601–1–051601–8.
- Winterstein, S., T. Ude, C. Cornell, P. Bjerager, & S. Haver (1993). Environmental parameters for extreme response: Inverse form with omission factors. In *Proc. 6th International Conference on Structural Safety and Reliability.*, Innsbruck, Austria, pp. 551–557. CRC Press, Taylor and Francis Group.