Impact of including interdependencies between multiple riverine flood defences on the economically optimal flood safety levels

Dupuits, E. J.C.; Klerk, W. J.; Schweckendiek, T.; de Bruijn, K. M.

DOI
10.1016/j.ress.2019.04.028

Publication date
2019

Document Version
Accepted author manuscript

Published in
Reliability Engineering and System Safety

Citation (APA)
Dupuits, E. J. C., Klerk, W. J., Schweckendiek, T., & de Bruijn, K. M. (2019). Impact of including interdependencies between multiple riverine flood defences on the economically optimal flood safety levels. *Reliability Engineering and System Safety*, *191*, Article 106475. https://doi.org/10.1016/j.ress.2019.04.028

Impact of including interdependencies between multiple riverine flood defences on the economically optimal flood safety levels

E.J.C. Dupuits[a], W.J. Klerk[a,b], T. Schweckendiek[a,b], K.M. de Bruijn[b]

[a] *Delft University of Technology, Faculty of Civil Engineering and Geosciences, P.O. Box 5048, 2600 GA Delft, Netherlands*
[b] *Deltares, P.O. Box 177, 2600 MH Delft, Netherlands*

Abstract

In risk analysis of riverine flood defence systems, sections of flood defences are often considered separately, herewith ignoring their interdependence, e.g. due to the hydraulic response following dike breaches in the system. In previous studies it has been found that such interdependence can have a significant influence on flood risk estimates and the spatial distribution. In this paper a method is proposed for the economic optimisation of riverine flood defence safety levels from a river system perspective. In order to deal with the computational challenge of integrating the hydraulic interactions in an economic optimisation, a surrogate model was developed. Despite the many simplifications, this model yields reasonably accurate results within acceptable time. The application of the model to a case study in the Netherlands has shown that taking into account interactions between flood defences has significant influence on optimal long term strategies for flood defences. The results suggest that accounting for interdependence in setting safety standards and reinforcement prioritisation yields a significant return on investment both in terms of lower investment cost and in terms of reduced risks.

Keywords: Economic optimisation, cost-benefit analysis, system reliability, flood risk, flood defences

1. Introduction

Settlements and industry along rivers are often protected against flooding by flood defences such as dikes and hydraulic structures. Multiple flood defences in the same river basin area can be considered as a system of riverine flood defences. In The Netherlands, flood defence systems are used to protect a major part of the country against flooding. The safety level of these systems is assessed periodically and policies are in place to meet the safety standard in 2050 for all primary flood defences.

A common way of determining how safe a flood defence system currently is, is by analysing for each flood defence separately how it performs under loading. In case of a riverine flood defence system, the loading could be a high river discharge. However, recent literature shows how multiple flood defences interact with each-other hydrodynamically as a system during an extreme event [e.g. 1, 2, 3, 4, 5]. These studies found significant differences in case a system as a whole was studied instead of as separate, independent elements.

The required safety of a flood defence can be (and, for the Netherlands, is) based on criteria for acceptable risk [6, 7]. One of the acceptable risk metrics used in the Netherlands is an economic cost-benefit analysis, which is used to determine the optimal protection level of a flood defence (e.g. see as early as [8]). In this paper we consider such an economic optimisation of protection levels. Other acceptable risk criteria such as societal risk and individual risk (life safety) are not considered here but are important nonetheless. Examples of other criteria for acceptable risk can be found in for example [9], as well as in the context of Integrated Flood Risk Management (IFRM, e.g. see [10]). However, in this paper only the economic acceptable risk by means of an economic optimisation is considered.

'Optimal' is defined in this context as where the total costs, which is the sum of investment costs and expected annual damages, is at a minimum. The Expected Annual Damages (EAD) of flood defences are defined in an economic optimisation as the expected loss in an arbitrary year, and can be found as a function of the annual probabilities of flooding and the expected damages due to flooding ([11]); an overview of this approach as followed in the Netherlands can be found in [12]. These EAD estimates change in time by effects such as economic growth and climate change. The results of following the general approach as described in [12] can be displayed as a series of (economically optimal) investments to adjust the safety level of a flood defence over time (Figure 1).

In an economic optimisation, the probability of flooding can be estimated using reliability analysis. In order to include interdependencies between flood defences in an economic optimisation, not only the reliability analysis need to be able to include these interdependencies, but the economic optimisation approach as well. In the economic optimisation methods mentioned in for example [12], interdependencies between multiple flood defences are not present. An economic optimisation method which is able

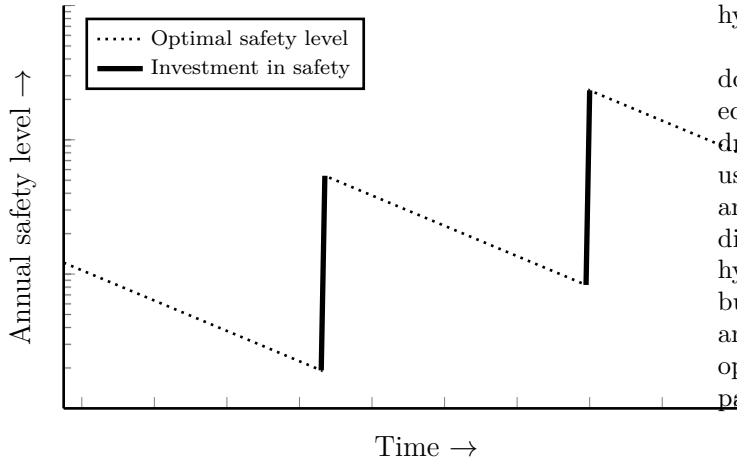

Figure 1: Schematic view of the result of an economic cost-benefit analysis for a flood defence (e.g. as described in [12]). A series of investments are needed to make sure that the defence remains at an optimal level from an economic point of view due to (over time) changing expected annual damages. Image reproduced from [13].

to handle these interdependencies (which was applied to a case study in the Netherlands) can be found in [14]. However, this method used pre-calculated EAD estimates and focused on an application for a coastal system. For riverine systems, this method needs to be evolved further as pre-calculating EAD estimates is not always feasible.

Generally speaking, each flood defence in a riverine system can be interdependent on any other flood defence in that same system by means of hydrodynamic interactions. A straightforward example of this is a river along which flood defences exist. If an upstream flood defence breaches, a part of the river discharge will flow through the breach, leading to less discharge downstream. Less discharge downstream means that the load on the downstream defences will be reduced. The fact that each flood defence in a riverine system can be interdependent on any other flood defence in that same system, makes a riverine flood defence system more complex than a coastal flood defence system. Furthermore, obtaining a single EAD estimate for a riverine flood defence system with hydrodynamic interactions can be computationally expensive: it can take hours [4] or even days [3]. In the context of an economic optimisation, where a large number of EAD estimates need to be evaluated, this can quickly become infeasible.

In this study a cost-benefit analysis is carried out for a riverine flood defence system with multiple interdependent flood defences, in a computational tractable manner. This economic cost-benefit analysis will then be used to compare the impact of including the effect of hydrodynamic interactions of multiple flood defences on the economically optimal investment scheme in a case study (compared to not including these interactions). The term interdependency, instead of the more generic (statistical) term dependency, is used here to express the dependent behaviour (i.e. hydraulic interactions) between different flood defences.

The work presented here builds upon previous work done in [15, 16, 13]. The general idea of performing an economic optimisation with EAD estimates based on hydrodynamic interactions between flood defences was also used [15], but with less types of hydrodynamic interactions and applied to a (different, simpler) coastal system. [16] discusses some fundamental impacts of including riverine hydrodynamic interactions on an economic optimisation, but does this using an analytical economic optimisation and hypothetical, small riverine systems. The economic optimisation method discussed in [13] is applied here as part of the case study.

2. Case description

The case study used in this research is based on the area in the Netherlands where the river Rhine enters the Netherlands, see also Figure 2. This area has been the subject of a number of recent studies regarding the impact of hydrodynamic interactions on flood risk estimates in the area (see for example [3, 4, 5]). Instead of using these existing models, we opted to develop a simplified model in order to reduce the model run time. This model is similar to the existing models with respect to the fact that it represents the impact of hydrodynamic interactions on EAD estimates. The underlying model is primarily based on the work in [4], but with simplified hydrodynamics (in order to reduce model run time). Section 3 contains further details regarding the model and how it is used in an economic optimisation. An overview of the case area is shown with two illustrations in Figure 3.

In Figure 3, the flood prone areas are 'D48', 'D49', 'D50' and 'D51'. All these areas can experience damage due to flooding (the areas are indicated with striped rectangles). Breach flows can occur due to breaches at nine locations, which represent potential breach locations and their impacts anywhere in the system. These breach locations are indicated with stars in Figure 3, while the arrows originating from the breaches represent the breach flows. The breach flows represent the hydrodynamic interactions in the model. Breaches can occur, depending on the breach location, external and/or internal. External breaches occur at the river-facing side of a breach location due to a local extreme river discharge. On the other hand, internal breaches occur due to a load at the polder-facing side due to water levels of an already flooded area.

The breach location 'GER', if it breaches, has two breach flows: one flow connects to the flood prone area D48 (10% of the breach volume Q_{GER}), while the second flow forms a shortcut to the river IJssel (90% of the breach volume Q_{GER}). Furthermore, the breach locations B49i and B50i can form internal shortcuts between the flood prone areas D49 & D50 and D50 & D51, respectively. The breach flows associated with the breach locations B49i & B50i are purposefully directed downstream, as the area

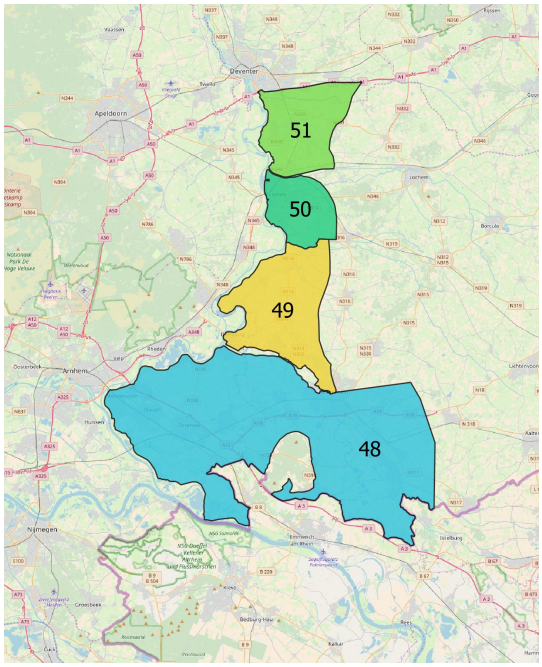

Figure 2: The actual area in the Netherlands upon which the case in this section is based. The numbered surfaces are dike ring areas, which are areas that are protected by flood defences (which are typically dikes). Map data © OpenStreetMap.

along the river IJssel is sloping downwards in the downstream direction of the river.

To provide a benchmark for results with hydrodynamic interactions, the model can be run without hydrodynamic interactions as well. This means that the hydrodynamic interactions are removed from the hydrodynamic model. Practically, 'no hydrodynamic interactions' has the following implications:

- The breach location GER will not form a shortcut from the Rijn to the IJssel;
- If a breach occurs, the downstream river discharge is unaltered;
- The internal breach locations B49i and B50i never breach, which means that internal shortcuts cannot be formed;
- The four breach locations GER, B481, B482 and B483 can all still influence the flood damage at D48 together.

3. Approach

3.1. General approach

In principle, an economic optimisation of a flood defence system is a cost-benefit analysis which attempts to minimize the Net Present Value (NPV) of the total cost, where the total cost is the sum of the Expected Annual Damage (EAD) estimates (Section 3.3) and the accompanying investment costs (Section 3.4); see also for example [8, 17]. In this case, the EAD estimates are determined by means of hydrodynamic simulations and impact assessments of (potential) flood events while taking into account the performance of the flood defence systems in place (Section 3.2). Practically, this means that the economic optimisation evaluates various system configurations (in terms of the reliability levels of the individual flood defence sections) in order to find an optimal investment scheme for the considered time period. A system configuration is defined here as a unique combination of flood defence levels. For example, in a flood defence system with two flood defences A and B, where both flood defences have five possible levels (labeled here as 1 to 5), a single system configuration (out of a total of $5^2 = 25$ possible configurations) would be flood defence A at level 1 and flood defence B at level 2.

Usually, an economic optimisation of flood defences requires repeated investments (see e.g. [18, 17]) as time dependent changes such a economic growth in a flood-prone area, reduced dike strength due to subsidence and increasing river discharges due to climatic changes increase the EAD. This means that, at some point in time, it will become economically attractive to (re)invest in reinforcing flood defences, which is incorporated in the optimisation routine through the Present Value (PV) of cash flows.

Figure 4 shows an overview of the approach. Starting from a current system configuration, the hydraulic interactions are implemented in hydrodynamic simulations and damage estimations (Section 3.2), which are used to estimate the EAD (Section 3.3). An optimisation algorithm then determines, based on a cost-benefit analysis, the optimal system configuration per time step for the considered time period (economic optimisation, Section 3.5). This collection of optimal system configurations in time can then be used to determine the accompanying optimal investment scheme.

3.2. Hydraulic simulations and damage estimations

If a flood defence breaches, flood damage due to the inflow of water can be expected. In order to estimate these (economic) flood damages, two primary elements need to be determined for a flood-prone area: the extent and severity of a flood, and the damage due to this flooding. The first can be simulated with a wide range of methods, for example with 1D models [e.g. as in 4], 2D models [e.g. as in 3], or flood cell storage methods [e.g. as in 1]. After estimating the extent and severity of the flood, typically in terms of flood depth and for some purposes also in terms of flow velocities, the information is combined with vulnerability characteristics (e.g. stage-damage curves) to obtain damage estimates. In The Netherlands, the HIS-SSM method can do this based on land use and damage functions [19, 4].

The hydraulic simulations for the case study are done with a (simplified) hydrodynamic model that describes the propagation of a (peak) discharge wave through the

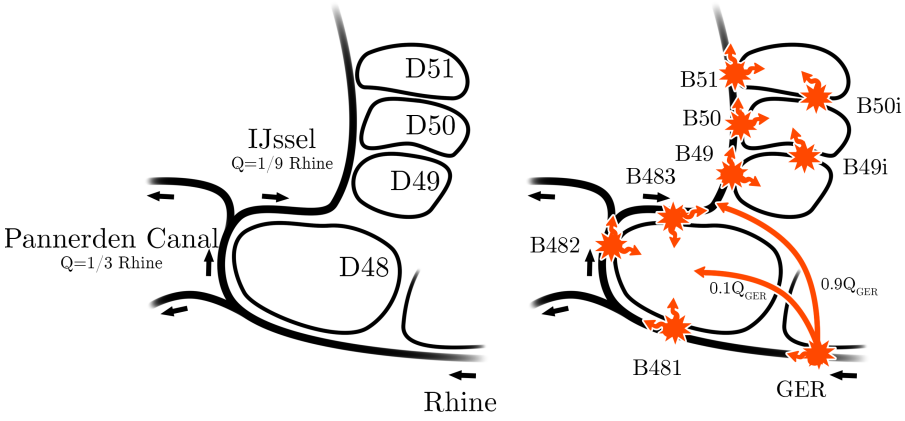

Figure 3: Overview of the simplified case study area. The illustration on the left shows the names and flow directions of the river branches in the study area, as well as the numbers of the areas prone to flooding. The illustration on the right shows the same area, but now with the locations & names of breaches and the resulting breach flows.

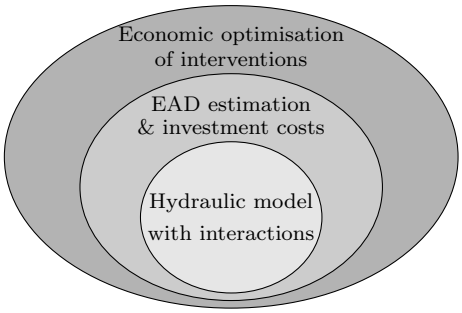

Figure 4: Overview of an economic optimisation with hydraulic interactions, as further described in Section 3.2 - 3.5.

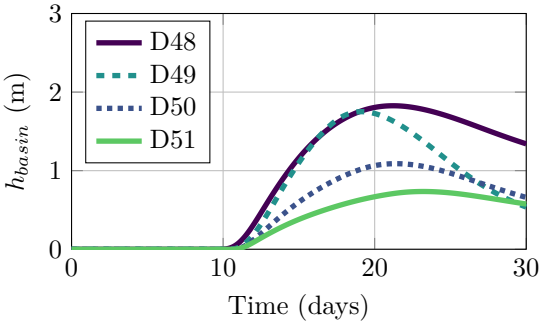

Figure 5: Example of inundation depths for the flood prone areas of Figure 3 for a single model run with a peak discharge of 16,000 m^3/s.

river branches. Once a breach occurs at a breach location, breach flows can occur toward flood-prone areas or form a shortcut between river branches (as described in Chapter 2). A more detailed description of the hydrodynamic model can be found in Appendix A. A single model run yields information regarding which flood defence has breached and what the inundation depths are at flood prone areas. Figure 5 shows an example of a time series of inundation depths with a time step size of two hours.

The damage (D) in the flood prone areas of the case study is assumed to follow a logarithmic relation that depends on the inundation depth (i.e. h_{basin}) and is shown in Eq. 1. The maximum damage (D_{max}) is reached at the inundation level d_{max}, as shown in Figure 6. The values for the maximum damage and the maximum inundation depth are shown in Table 1. The values for the maximum damage (D_{max}) are based on data listed in [17] which represent a monetary valuation of material and non-material loss in case of a flood. Furthermore, the maximum damage is assumed to increase over time (t, in years). In accordance with [17], the annual economic growth rate (γ) is set to 0.02.

$$D\left(h_{basin}, t\right) = \begin{cases} D_{max} \frac{\ln(1+h_{basin})}{\ln(1+d_{max})} e^{\gamma t} & \text{if } h_{basin} \leq d_{max} \\ D_{max} e^{\gamma t} & \text{if } h_{basin} > d_{max} \end{cases} \quad (1)$$

Flood prone area	D_{max} (10^6)	$d_{max}(m)$
D48	7046	5.0
D49	82	5.0
D50	2119	5.0
D51	57	5.0

Table 1: D_{max} and d_{max} values for the flood prone areas of Figure 3. The D_{max} values are based on data listed in [17].

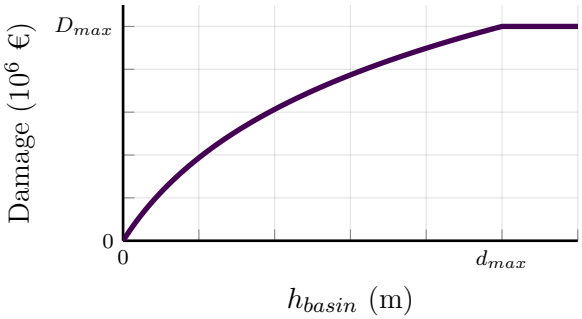

Figure 6: Damage at a flood prone area grows logarithmically to a maximum damage value D_{max}, which coincides with an inundation level at d_{max}.

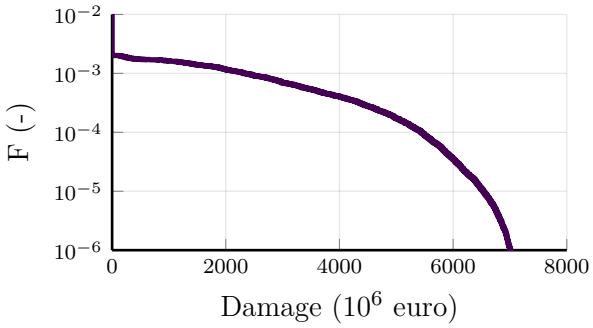

Figure 7: Example of a flood loss curve for D48 with hydrodynamic interactions, which plots the annual exceedance probability F versus the damage. In this example, all breach locations were set to an annual failure probability of 1/1000.

3.3. EAD estimation

On an abstract level, the likelihood of flooding depends on the ratio between the strength and the load of the flood defences. This is shown in the reliability equation Z in Eq. 2.

$$Z = Strength - Load \tag{2}$$

In this paper the load is the water level at the defence and the strength is the critical height of the flood defence, which is uncertain due to the nature of most flood defence failure mechanisms (e.g. shear resistance of the soil in and under a dike). The evaluation of the limit state function results in an estimation of the probability of failure (of that breach location).

Eq. 2 is evaluated for each breach location in the case study. Various methods are available to estimate the probability of failure with this equation, of which Monte Carlo simulation is a frequently used method; regarding (Dutch) riverine flood defence systems with hydrodynamic interactions, see for example [3] and [5]. We use a similar Monte Carlo method with Importance Sampling as mentioned in [20]. This method (including contributing distributions) is further specified for the case study in Appendix A.3.

The Expected Annual Damage (EAD) estimates in the case study are determined with the Monte Carlo method as well. These estimates are calculated by determining (for each Monte Carlo sample) the maximum water depth in a flood prone area and the accompanying flood damage (i.e. Section 3.2). By incorporating the likelihood of the Monte Carlo sample (and hence the likelihood of the flood damage), a flood loss curve can be constructed as shown in Figure 7. The area under this flood loss curve represents the EAD estimate for that specific flood prone area. By using this approach, the EAD estimate can be influenced by any breached breach location ((or multiple breach locations) that leads to a damage at the flood prone area. This is essential in order to determine the effects of hydrodynamic interactions on the EAD estimates.

3.4. Investment costs

The investment costs represent the cost of increasing the safety level of a flood defence. These costs can be determined based on actual design studies for specific locations, but can also be approached by more general relations. Specifically for dikes, a number of relations were discussed in, for example, [17]. The exponential, dike-height dependent relation proposed in [17] is shown in Eq. 3. In Eq. 3, u is defined as the height increase from height h_1 to height h_2, C_f is the fixed investment cost, C_v is the variable investment cost, and λ is an exponential scaling factor. This equation will be used in the case study of this paper.

$$I\left(u, h_2\right) = \begin{cases} 0 & \text{if } u = 0 \\ \left(C_f + C_v u\right) e^{\lambda h_2} & \text{if } u > 0 \end{cases} \tag{3}$$

Investment costs and time dependent parameters for the case study are associated with each breach location, and are shown in Table 2. Similar to the maximum damage values (Section 3.2), the values in Table 2 are based on data found in [17]. The parameters C_f, C_v and λ are to be used with the exponential investment relation (Eq. 3). Parameter η has been interpreted as a proxy for degradation of strength over time in [17]. Therefore, it was implemented as a reduction on the mean critical height of the associated breach location. Furthermore, the yearly discount rate δ is set to 0.04 (in accordance with [17]).

A slight modification from [17] is made for D48, as it has only a single value and contains three breach locations in this case study (see Figure 3). Therefore the values for investment cost related parameters C_f and C_v are distributed equally over the breach locations B481, B482 and B483. Furthermore, the breach location GER has been given the same investment and time dependent characteristics as B481 (and therefore B482 & B483).

Because of the time dependent parameters, the number of possible decisions to be evaluated in the economic optimisation (Section 3.5) is influenced not only by the number of system configurations, but also by the considered reinforcement moments in time. Specifically for the number of EAD calculations, the considered reinforcement moments in time increase the number of EAD calculations as the EAD estimates differ each year.

Breach	C_f (10^6)	C_v (10^6)	λ (-)	η $(m/year)$
GER	11.9	47.7	0.63	0.00496
B481	11.9	47.7	0.63	0.00496
B482	11.9	47.7	0.63	0.00496
B483	11.9	47.7	0.63	0.00496
B49	20.0	80.0	0.46	0.00304
B50	8.13	33.0	0.00	0.00320
B51	15.0	60.0	0.71	0.00294

Table 2: Investment and time dependent parameter values for the breach locations of Figure 3. The values mentioned for C_f, C_v and λ are input for the investment function discussed in Section 3.4 and based on data in [17]. η denotes the degradation of strength over time.

3.5. Optimisation routine

Finding the economically optimal investment scheme can be done analytically for simple systems (e.g. see [17]). For larger (or more complex) flood defence systems numerical methods are more convenient. A recent numerical method that is capable of optimising flood defence systems with multiple interdependent defences can be found in [14] and [21]. The study in [13] uses a similar approach as in [14] to optimise flood defence systems with multiple interdependent defences, but attempts to reduce the number of EAD calculations required by the numerical optimisation method.

The optimisation process can be conveniently visualized with the help of graphs [22]. This is shown conceptually in Figure 8. In this plot, each dot represents a system configuration at a certain moment in time. The lines between the dots represent a change in system configuration (if the dot later in time is at a higher position on the y-axis). The lines are given weights based on the sum of investment costs and EAD between, for example, t_0 and t_i, in order to obtain the optimal system configuration. Even if the system configuration does not change, the EAD can change due to temporal changes (see also Section 3.3 & 3.4).

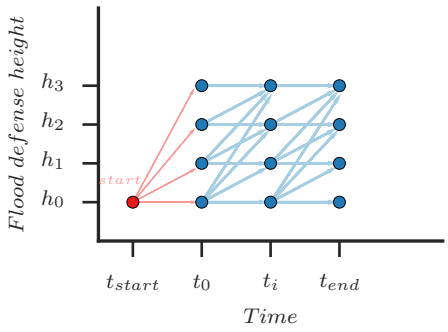

Figure 8: Conceptual image of a graph for a single flood defence, which helps to visualize the optimisation process. Image reproduced from [13].

Using the numerical optimisation method of [13], the number of EAD calculations is reduced by only executing the calculation of the necessary estimates which are used by the optimisation routine (i.e. 'lazy evaluation'). This method was shown to reduce the number of required EAD calculations by not having to calculate EAD estimates at lower safety levels in the distant future. Based on the examples in [13], the expected reduction is roughly a factor two.

For the case study, five different levels are considered per defence, in accordance with levels currently considered in Dutch safety standards (e.g. see [7]) and are associated with the failure probabilities per breach location. These five levels are called L1 to L5 and are related to the current ($t = 0$) annual failure probabilities of 1/300, 1/1000, 1/3000, 1/10,000 and 1/30,000 for each breach location separately *without* interactions. The number of system configurations for Figure 3, with the three internal breach locations linked to the levels of the external breach locations (see Appendix A.3), is equal to 5^7 or $78,125$.

The total time horizon used in the economic optimisations for the case study spans 300 years in the future (similar to the choice in e.g. [23]), with 58 moments marked as potential reinforcement times. These 58 moments are concentrated mostly in the near future. This is because decisions in the near future are considered as more important than decisions in the distant future (i.e. weigh more heavily on the total cost estimate due to discounting). The first 20 years have a possible decision each year (including one at $t = 0$), while the next 80 years have a possible decision every five years, while the remaining 200 years have a possible decision every ten years.

3.6. Computational efficiency

Numerical modelling of hydrodynamic interactions, especially in a Monte Carlo setting, can be computationally expensive, see for example [3]. This computational burden will be amplified in the context of an economic optimisation, which investigates multiple system configurations (see also Section 3.5). Each system configuration needs an EAD estimate; which means multiple Monte Carlo simulations. The number of system configurations is dependant on the underlying system, but can easily reach hundreds or thousands system configurations.

Without hydrodynamic interactions (see also Section 2), the economic optimisation can be done independently for each flood prone area. Practically, this means that for D48 there are 625 possible system configurations per time step (5^4), while for the other flood prone areas (D49, D50 and D51) there are only five possible system configurations per time step.

If hydrodynamic interactions are included, the number of potential system configurations increases to 78,125 per time step (Section 3.5). This results in more than 23 million potential EAD calculations in a period of 300 years. If only 50% of these 23 million EAD calculations actually need to be computed and each EAD calculation takes about one second, the resulting computational time would be approximately 68 days on a single CPU core. If these

calculations can be distributed with perfect efficiency over multiple cores, 100 cores would be done in less than a day. However, during this study a computer cluster was not available which meant that computing the EAD estimates were considered a high computational burden. This burden would be even larger if a more complex hydrodynamic model was used (e.g. in [4] it took hours to compute a single EAD estimate).

Because of this high computational burden, it was first checked whether or not including hydrodynamic interactions leads to significantly different results. To that end, Spearman's correlation coefficients were calculated between the input (critical heights of the breach locations) and output (EAD per flood prone area) for all 78,125 system configurations at $t = 0$ with hydrodynamic interactions in Table 3.

Breach location	D48	D49	D50	D51
GER	0.01	-0.54	-0.49	-0.39
B481	-0.65	0.05	0.05	0.05
B482	-0.47	0.09	0.09	0.08
B483	-0.31	0.16	0.13	0.12
B49	0.00	-0.75	0.07	0.17
B50	0.00	0.00	-0.78	-0.02
B51	0.00	0.00	0.00	-0.82

Table 3: Spearman's correlation coefficients for the critical heights of breach locations (rows), versus the EAD of the flood prone areas (columns).

Table 3 shows that the breach locations upstream of a flood prone area (see also Figure 3) have a significant correlation with the EAD in that flood prone area. These correlation coefficients show that the hydrodynamic interactions cannot be disregarded a-priori and that an economic optimisation with hydrodynamic interactions will most likely produce different results than an economic optimisation without hydrodynamic interactions. Therefore, surrogate modelling was applied to reduce the calculation time.

Surrogate modelling is an approximation method where computationally expensive models are replaced with more efficient surrogates; a review of surrogate modelling can be found in [24]. In this review, two types of surrogate models are distinguished: physically based models with lower numerical complexity (e.g. going from a 2D model to a 1D model), and response surface surrogates (e.g. fitting a polynomial function through results obtained from a 2D model). Surrogate modelling can therefore help in relieving the computational burden, at the cost of returning an approximation of the results of a model with higher numerical complexity rather than running the actual model over and over again. Circumventing computationally expensive Monte Carlo simulations has been done before (e.g. see the list of applications in [24]). Either way, we consider it as an important part of fulfilling our aim because of the necessary computational savings.

In this study, an Artificial Neural Network (ANN) was chosen to approximate the EAD calculations, as the case study already uses a hydrodynamic model with a low numerical complexity. An ANN was used as the response of the underlying data was unknown; neural nets are perceived to have a high flexibility towards emulating the response of underlying data. A more detailed description of the basic structure and training of neural network can be found in for example [25].

In order to train the neural network, all 78,125 possible system configurations at time step $t = 0$ were calculated. This data size could be calculated in a couple of hours because of the relatively simple hydrodynamic relations in the case study. Within the training data set, we did observe that the response type seems to be non-linear, as shown in the plots of Figure 9. This justifies the usage of a neural net instead of using, for example, a simpler linear regression analysis. Further details regarding the implementation and performance of the neural networks are described in Appendix B. In Appendix B, the correlation coefficients of Table 3 are well approximated using the surrogate model. With the help of the surrogate models, the economic optimisation takes minutes instead of weeks.

4. Results

In this section, we apply the approach of Section 3 to the case as described in Section 2 to estimate the impact of including hydrodynamic interactions in a riverine flood defence system with multiple dike sections. This is done by comparing EAD estimates and optimal system configurations with and without accounting for hydrodynamic interactions. Based on literature, a significant difference is expected to be found in the flooding probabilities and associated EAD estimates for simulations with and without interactions; see for example [3, 5, 4, 16].

In order to check whether or not the model with interactions yields different results than the model without interactions, a qualitative check is made for area D48, with all breach locations set to an annual failure probability of 1/1000. The cumulative probability distribution of flood damage (flood loss curve) for both models is shown in Figure 10. This figure shows a clear difference for flood prone area D48 where between the two model runs from about 3 billion euros up until the maximum flood damage. The tails of the flood loss curves converge because in both curves the maximum damage in D48 is reached (Table 1).

4.1. Optimisation without hydrodynamic interactions

To provide a benchmark for an economic optimisation with hydrodynamic interactions, also an economic optimisation without hydrodynamic interactions was done. The resulting optimal investment schemes for the seven breach locations are shown in Figure 11 & 12. The expected behavior per flood defence is a gradual increase of the level over time. This is explained by economic growth (i.e. larger potential damages) and a higher probability of a

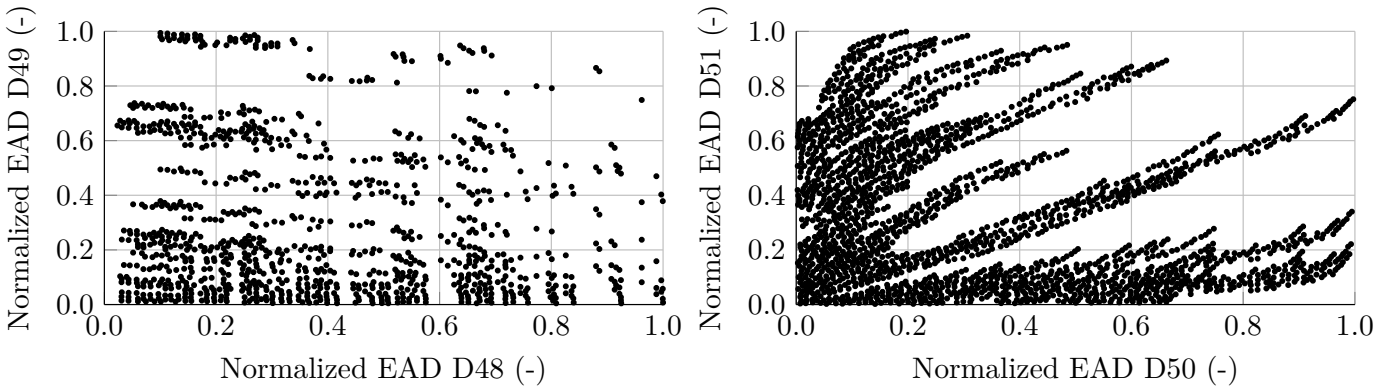

Figure 9: Normalized EAD scatter plots for identical system configurations at $t = 0$; all possible 78,125 system configurations are used. Two possible combinations of flood prone areas are shown, D48 versus D49 and D50 versus D51. The response seems to be weak non-nonlinear in D48 versus D49, while there seems to be a strong non-linear response in D50 versus D51. If the areas were hydrodynamically independent, the plots would be expected to be an evenly distributed field of dots.

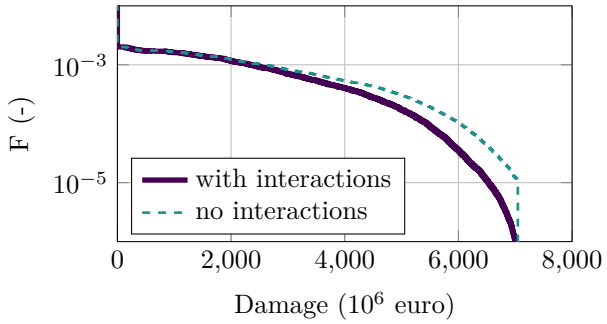

Figure 10: Example of flood loss curve for D48 with and without hydrodynamic interactions, which plots the annual exceedance probability F versus the damage. All breach locations were set to an annual failure probability of 1/1000.

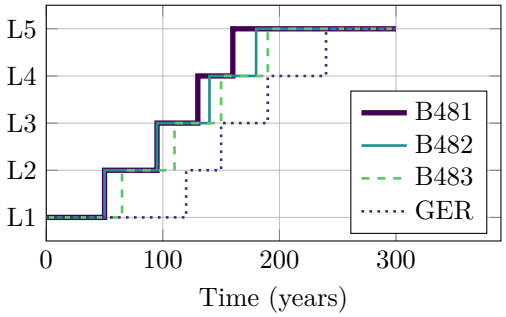

Figure 11: Optimal investment schemes assuming no hydrodynamic interactions for the four breach locations of D48 (GER, B481, B482 and B483).

flood defence failing (e.g. due to flood defence degradation or climate change). Breach location B50 in Figure 12 is the only location that does not show a gradual increase in level. Apparently, flood prone area D50 is attractive for investment (i.e. investment costs are relatively low regarding the achievable EAD reduction), and the initial level L1 is too low. The maximum level of L5 is already reached around year 60, which indicates that the chosen safety levels are possibly limiting the economic optimisation for B50; in the following years a higher level than L5 might, economically, be a better optimal choice. Other locations hit their maximum level later: around year 200, or not at all (i.e. B51).

The flood prone areas can be optimised separately as hydrodynamic interactions are not included in these optimisations, which reduces the number of EAD calculations. Because the EAD over a period of multiple years is approximated as the sum of EAD estimates for that period, the maximum number of EAD calculations for D48 is approximately the number of years plus one (301) times the number of system configurations (625), which equals 188,125 potential EAD calculations. By using the lazy

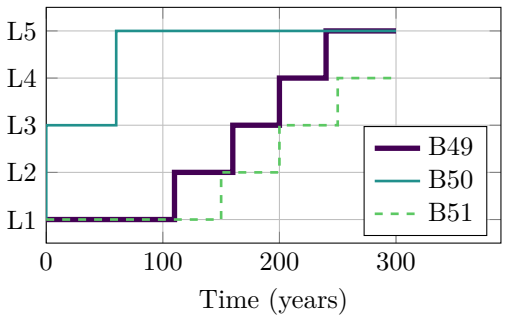

Figure 12: Optimal investment schemes assuming no hydrodynamic interactions for the three breach locations for D49, D50 and D51 (B49, B50 and B51).

evaluation feature of the used economic optimisation routine (see Section 3.5), 52% of these EAD calculations were actually executed.

4.2. Optimisation with hydrodynamic interaction

With the help of the surrogate models as described in Section 3.6, the economic optimisation takes minutes instead of weeks. Even though the surrogate model reduces the computational impact of acquiring EAD estimates considerably, this was further reduced by the lazy execution of the used optimisation routine (see Section 3.5): about 60% of all potential EAD calculations was actually used.

The economic optimisation with hydrodynamic interactions is carried out in two variants: one with breach location GER at a constant level (L1), and one with breach location GER as any other breach location, free to be optimised. This was done as the breach location GER in the case study represents a breach location in Germany. Therefore, this breach location may not, or cannot, be subjected to an investment strategy as desired by the Dutch part of the system.

Figure 13 shows the temporal development of the total EAD on a system level for the two strategies and the case without interactions. This shows that the development of the EAD under the assumption of 'no interactions' and 'with interactions and GER free' is roughly the same.

For 'GER fixed', the EAD clearly deviates after about 150 years. It is likely that, because GER is fixed at a low level, the economic optimisation routine is not able to 'control' the growth of the EAD elsewhere as investments here are no longer cost-efficient. Also the growth in EAD might not be contained because the highest available level for breach locations (L5) is simply not high enough.

Figure 14 shows the development of the annual flooding probabilities over time for the optimal investment schemes found with and without hydrodynamic interactions, using the model *with* hydrodynamic interactions. Especially the development of D50 using the investment scheme without interactions can be seen as too conservative (i.e. these flooding probabilities are significantly lower than the investment scheme with interactions). This is further supported by Figure 15, which shows the cumulative discounted investment costs of the three investment schemes: the investment costs for the scheme based on the case without interactions are significantly higher than the investment costs for both schemes determined with interactions.

The optimal investment schemes for the optimisation without hydrodynamic interactions, with interactions and with interactions but with GER fixed are shown in Figures 16, 17 & 18, respectively. When compared to the situation without interactions, it is clear that both optimisations with interactions have a significantly different investment strategy. In the first 100 years, investments in B481 and B482 are only slightly delayed when compared to the optimisation without interactions, as the discharge reduction from upstream breaches is small compared to

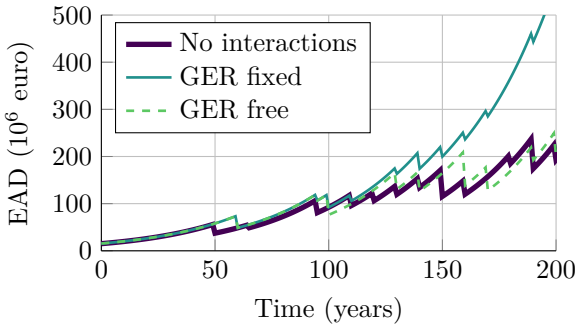

Figure 13: System EAD (i.e. summed EAD of all flood prone areas) over time for the three optimal investment paths as done in this section: one with no interactions (Section 4.1) and two with interactions (one with breach location GER fixed at level L1 and one with GER free to be optimised).

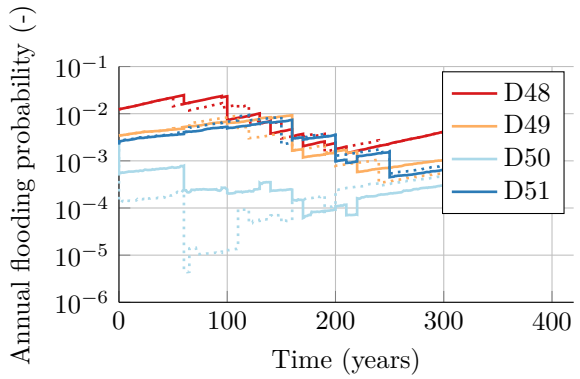

Figure 14: Annual flooding probability over time using the optimal investment scheme which was determined *with* interactions (solid lines, 'GER free') and *without* interaction (dotted lines). The flooding probabilities are in both cases calculated *with* interactions.

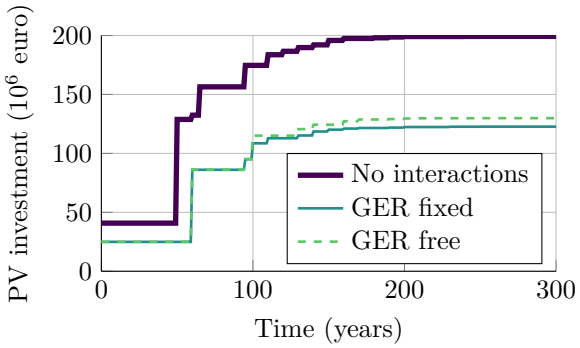

Figure 15: Cumulative sum (over time) of the present value of the system investment costs using the optimal investment schemes which were determined in this section.

the Rijn/Nederrijn discharge. For the locations at the IJssel (B483, B49, B50 and B51), investments are delayed significantly and (initial) investments are reduced as well, as the effect of upstream breaches and shortcutting is larger due to the smaller discharge capacity of the IJssel. The fact that for 'GER free' investment at these three locations is reduced further compared to 'GER fixed' highlights the importance of reinforcing dike section GER in order to manage downstream flood risks: a low level for dike section GER leads to higher risk at B49, B50 and B51 caused by higher water levels due to shortcutting.

Furthermore, Figure 18 with breach location GER fixed at L1 shows a peculiar large jump in safety level for B51 around year 230. This can be explained by the fact that all other options except increasing B49 from L4 to L5 are exhausted. However, as B49 also reduces the flood probability of B50 (where the potential damage is much higher) that investment might only increase the risk further, meaning that investment in B51 is the only option to (slightly) mitigate the exponential growth in EAD shown in Figure 13.

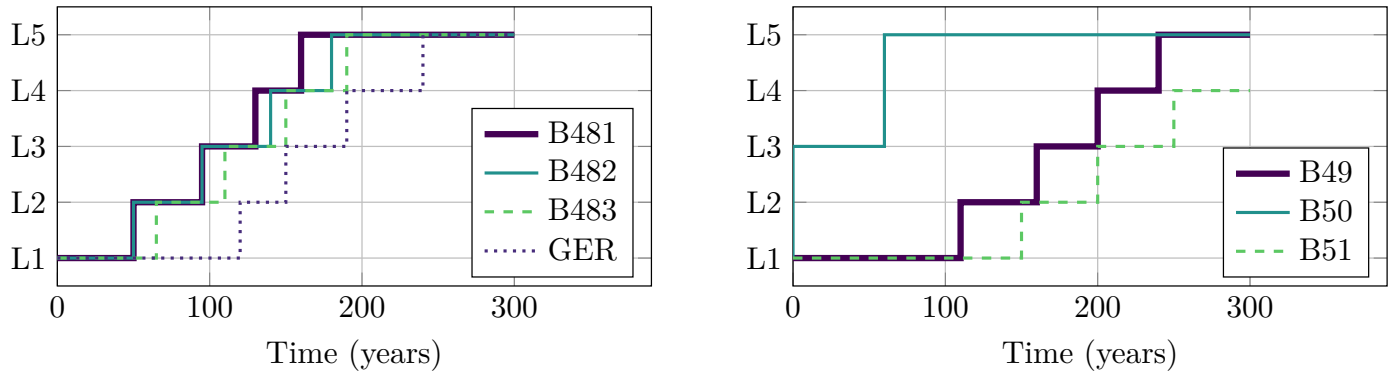

Figure 16: Investment schemes for the seven breach locations that follow from the economic optimisation without hydrodynamic interactions (Section 4.1, copies of Figure 11 & 12).

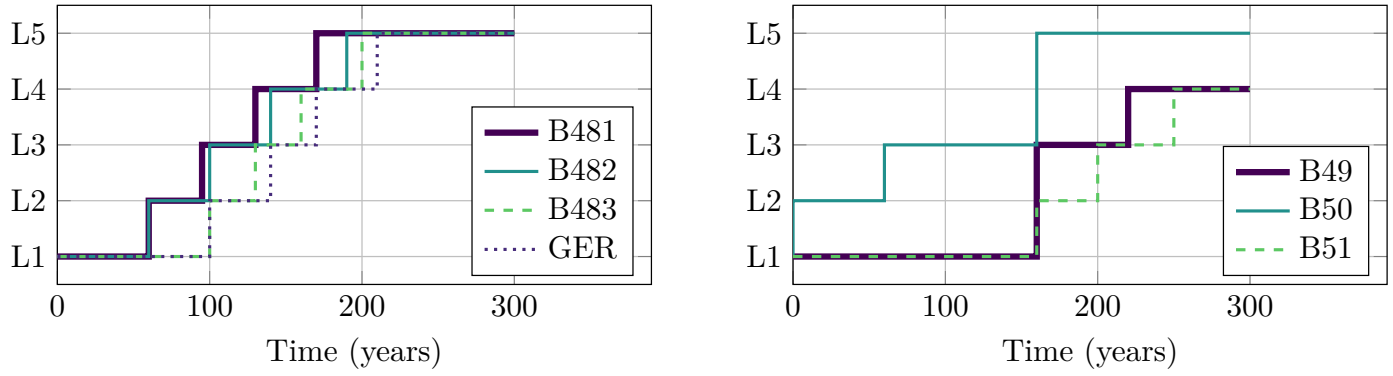

Figure 17: Investment schemes for the seven breach locations for an economic optimisation with hydrodynamic interactions and breach location GER free to be optimised as well.

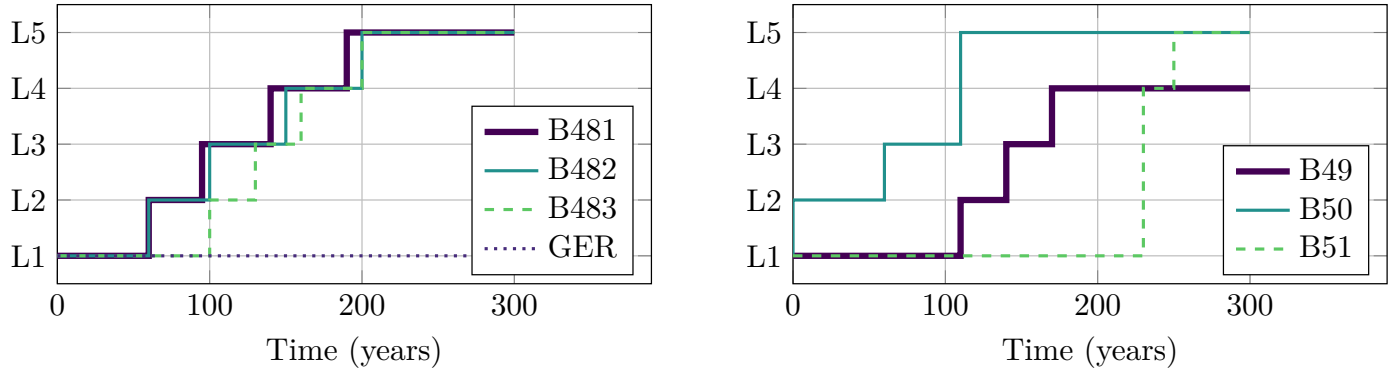

Figure 18: Investment schemes for the seven breach locations for an economic optimisation with hydrodynamic interactions and breach location GER fixed at its L1 level.

5. Discussion

In this paper we have presented a case where interdependence between riverine flood defences has been taken into account in the context of an economic cost-benefit optimisation. Whereas past studies on this topic, such as the ones by [4], [3] and [2] only considered the phenomenon itself, here we have added the perspective of optimal decision making and investment patterns. Previous studies on optimal investments for riverine flood defences (e.g. [12]) also did not consider these interdependencies.

As the computational complexity of these calculations is large it was necessary to simplify some aspects of the model, such as the hydrodynamic model. Despite the simplification of the hydrodynamic model the general behaviour of the system is still very much in line with the model used by [4] who used a calibrated quasi-2D model. Nevertheless it could be valuable to use a more complex calibrated model in an optimisation context as it will also improve assumptions such as the constant discharge fractions of the various river branches (see Figure A.21).

Incorporating such a more complex model in this context would however significantly increase computation times of the optimisation model. A possible avenue to mitigate this could be to calibrate a simple model such as described in Section 2 using the complex model. This would still improve computation time, and would in essence be a replacement of a model with high numerical complexity with a surrogate model of lower numerical complexity (see also Section 3.6). At the same time, it might be more efficient to train a response surface surrogate model directly on the output of a more complex model, without the intermediate step of a model with lower numerical complexity.

For more complex models, defining a proper training data set poses some additional issues. In Section 4.2, all system configurations at a single time step are used as the training data set. Especially for more complex models and larger systems calculating all system configurations might become infeasible. To that end, using smaller training data sets might be feasible as well (e.g. see [26]), or Latin hypercube sampling could be used to optimise the size of the training dataset while still achieving a good fit; see also [24].

Even though the framework applied provided reasonably accurate results in acceptable calculation times, modelling choices made in this study could be altered for other cases. This involves the choice of surrogate model type and which parts of the calculation to include in the surrogate model. Such choices should always be based on a consideration of required accuracy versus computational tractability.

It is found from the results that whether or not interdependencies are included in the model makes a large difference for the investment pattern, especially for the more downstream flood defence sections. This is in line with the findings by [4] on the same area, where it was found that the influence of interdependencies on risk levels was larger for more downstream locations. Up to now these influences have not been included in quantitative analysis of optimal safety standards [12] or investment patterns. The most prominent reason for this is the computational complexity, however techniques such as the neural networks used here are promising in tackling such complex optimisation problems [24].

Where we studied optimal investment patterns in this study, another potentially promising application of the modelling approach is the prioritisation of reinforcement measures under budget constraints. As shown in this study, the investment costs can change significantly if hydrodynamic interactions are included. For example, if the estimate of a flooding probability decreases by including interdependencies, a reinforcement (investment) of that flood defence section can be postponed in favour of other more urgent defences. This is particularly relevant for large flood defence infrastructure investment programs such as the Flood Protection Program in the Netherlands.

The case study considers an area close to the border with Germany. In this area dike breaches in Germany can cause significant damage in the Netherlands, meaning that the optimisation of the German flood defences should also be considered in the optimisation for the Dutch system. In this paper we have therefore included the cases 'GER free' and 'GER fixed'. From a comparison of these cases we found that especially on the long term the EAD skyrocketed, as the EAD is dominated by breaches in Germany. This illustrates the importance of looking past administrative borders in order to achieve appropriate flood risk management strategies.

In the case study in this paper we have shown that interdependencies can have a large influence on flood risk management strategies. In order to fully exploit the model set-up outlined in this paper for decision making on safety standards, prioritisation and cross-border risk analysis it is specifically important that there is a shared trust towards the underlying models. This interest potentially conflicts with simplifying the model in order to keep the computational burden of these calculations in check. Therefore it is important that future developments focus on connecting the outlined approach to calibrated models, for which some suggestions have been presented in this section.

6. Conclusions

In this study we presented a modelling framework that enables an economic cost-benefit analysis of a riverine flood defence system with multiple interdependent flood defences. This economic cost-benefit analysis was then used to compare the impact of including the effect of multiple interdependent flood defences and their hydrodynamic interactions on the economically optimal investment scheme in a case study (versus not including these interactions).

Using a simple hydrodynamic model in a Monte Carlo simulation with Importance Sampling enabled economic optimisation with limited or no hydrodynamic interactions

(as described in Section 4.1). However, to be able to take into account the increasing number of system configurations when accounting for hydrodynamic interactions, a neural network was used as a surrogate model. We have shown that for the case study in this paper, the used surrogate model provides a reasonable approximation for the simple hydrodynamic model whilst significantly reducing the computational burden. For the case study, significant differences were found both in terms of the timing of the optimal investments, as well as the magnitude of the required investments.

We have shown that for efficient flood risk management strategies in the case study area interdependencies have to be taken into account as there are significant differences in investment patterns. Therefore relaxing the assumptions of independence will likely lead to more optimal investment strategies and thus more cost-effective protection of flood prone areas. As there are several options for coping with the typically significant computational burden of including hydrodynamic interactions we expect that the research costs towards including hydrodynamic interactions will be significantly outweighed by the potential cost-savings that can be achieved by having a better, more optimal, investment scheme.

Acknowledgements

We are grateful for the financial support of the Netherlands Organization for Scientific Research, Domain Applied and Engineering Sciences (NWO-TTW), which is partly funded by the Dutch Ministry of Economic Affairs.

[1] S. Vorogushyn, B. Merz, K.-E. Lindenschmidt, H. Apel, A new methodology for flood hazard assessment considering dike breaches, Water Resources Research 46 (8). doi:10.1029/2009WR008475.
URL http://dx.doi.org/10.1029/2009WR008475

[2] S. Vorogushyn, K.-E. Lindenschmidt, H. Kreibich, H. Apel, B. Merz, Analysis of a detention basin impact on dike failure probabilities and flood risk for a channel-dike-floodplain system along the river Elbe, Germany, Journal of Hydrology 436-437 (2012) 120–131. doi:10.1016/j.jhydrol.2012.03.006.
URL http://www.sciencedirect.com/science/article/pii/S0022169412001928

[3] W. Courage, T. Vrouwenvelder, T. van Mierlo, T. Schweckendiek, System behaviour in flood risk calculations, Georisk: Assessment and Management of Risk for Engineered Systems and Geohazards 7 (2) (2013) 62–76. doi:10.1080/17499518.2013.790732.
URL http://dx.doi.org/10.1080/17499518.2013.790732

[4] W. Klerk, M. Kok, K. de Bruijn, S. Jonkman, P. van Overloop, Influence of load interdependencies of flood defences on probabilities and risks at the Bovenrijn/IJssel area, The Netherlands, in: Proceeding of the 6th international conference on flood management - ICFM6, 1-13., Brazilian Water Resources Association and Acquacon Consultoria, 2014.

[5] K. M. De Bruijn, F. L. M. Diermanse, J. V. L. Beckers, An advanced method for flood risk analysis in river deltas, applied to societal flood fatality risks in the Netherlands, Natural Hazards and Earth System Sciences Discussions 2 (2) (2014) 1637–1670. doi:10.5194/nhessd-2-1637-2014.

[6] J. Vrijling, W. van Hengel, R. Houben, Acceptable risk as a basis for design, Reliability Engineering & System Safety 59 (1) (1998) 141–150. doi:10.1016/S0951-8320(97)00135-X.
URL http://linkinghub.elsevier.com/retrieve/pii/S095183209700135X

[7] M. Kok, R. Jongejan, M. Nieuwjaar, I. Tanczos, Fundamentals of Flood Protection, Tech. rep., Expertise Netwerk Waterveiligheid (ENW) (2017).

[8] D. Van Dantzig, Economic Decision Problems for Flood Prevention, Econometrica 24 (3) (1956) 276–287.

[9] S. Jonkman, P. van Gelder, J. Vrijling, An overview of quantitative risk measures for loss of life and economic damage, Journal of Hazardous Materials 99 (1) (2003) 1–30. doi:10.1016/S0304-3894(02)00283-2.

[10] P. Samuels, M. Morris, P. Sayers, J. Creutin, A. Kortenhaus, F. Klijn, E. Mosselman, A. Van Os, J. Schanze, A framework for integrated flood risk management, 2010.

[11] B. Gouldby, P. Samuels, Language of Risk. Project definitions, FloodSite.

[12] J. Kind, Economically efficient flood protection standards for the Netherlands, Journal of Flood Risk Management 7 (2) (2014) 103–117. doi:10.1111/jfr3.12026.
URL http://doi.wiley.com/10.1111/jfr3.12026

[13] E. J. C. Dupuits, F. L. M. Diermanse, M. Kok, Economically optimal safety targets for interdependent flood defences in a graph-based approach with an efficient evaluation of expected annual damage estimates, Natural Hazards and Earth System Sciences 17 (11) (2017) 1893–1906. doi:10.5194/nhess-17-1893-2017.
URL https://www.nat-hazards-earth-syst-sci.net/17/1893/2017/

[14] P. Zwaneveld, G. Verweij, Economisch optimale waterveiligheid in het IJsselmeergebied, Tech. Rep. 10, CPB, The Hague (2014).

[15] E. Dupuits, T. Schweckendiek, M. Kok, Economic Optimization of Coastal Flood Defense Systems, Reliability Engineering & System Safety 159.

[16] E. Dupuits, K. de Bruijn, F. Diermanse, M. Kok, Economically optimal safety targets for riverine flood defence systems, E3S Web Conf. 7 (2016) 20004. doi:10.1051/e3sconf/20160720004.
URL https://doi.org/10.1051/e3sconf/20160720004

[17] C. Eijgenraam, R. Brekelmans, D. den Hertog, K. Roos, Optimal Strategies for Flood Prevention, Management Science-doi:10.1287/mnsc.2015.2395.
URL http://dx.doi.org/10.1287/mnsc.2015.2395

[18] J. Vrijling, I. van Beurden, Sealevel rise: a probabilistic design problem, Coastal Engineering Proceedings (1990) 1160–1171.
URL http://journals.tdl.org/icce/index.php/icce/article/viewArt

[19] M. Kok, Standaardmethode 2004 : schade en slachtoffers als gevolg van overstromingen, Ministerie van Verkeer en Waterstaat, Rijkswaterstaat, Den Haag, 2004.
URL http://edepot.wur.nl/78084

[20] F. L. M. Diermanse, K. M. Bruijn, J. V. L. Beckers, N. L. Kramer, Importance sampling for efficient modelling of hydraulic loads in the Rhine–Meuse delta, Stochastic Environmental Research and Risk Assessment 29 (3) (2015) 637–652. doi:10.1007/s00477-014-0921-4.
URL http://dx.doi.org/10.1007/s00477-014-0921-4

[21] B. Yüceoglu, Branch-and-cut algorithms for graph problems, Ph.D. thesis, Maastricht University (2015).

[22] P. J. Zwaneveld, G. Verweij, Safe Dike Heights at Minimal Costs, Tech. rep., CPB, The Hague (2014).

[23] R. Brekelmans, D. Den Hertog, K. Roos, C. Eijgenraam, Safe Dike Heights at Minimal Costs: The Nonhomogeneous Case, Operations Research 60 (2012) 1342–1355. doi:10.1287/opre.1110.1028.

[24] S. Razavi, B. A. Tolson, D. H. Burn, Review of surrogate modeling in water resources, Water Resources Research 48 (7) (2012) n/a—-n/a. doi:10.1029/2011WR011527.
URL http://dx.doi.org/10.1029/2011WR011527

[25] A. A. Chojaczyk, A. P. Teixeira, L. C. Neves, J. B. Cardoso, C. Guedes Soares, Review and application of Artificial Neural Networks models in reliability analysis of steel structures, Structural Safety 52 (PA) (2015) 78–89. doi:10.1016/j.strusafe.2014.09.002.
URL http://linkinghub.elsevier.com/retrieve/pii/S016747301400085

[26] J. P. Aguilar-López, J. J. Warmink, R. M. J. Schielen, S. Hulscher, Piping erosion safety assessment of

flood defences founded over sewer pipes, European Journal of Environmental and Civil Engineering (2016) 1–29doi:10.1080/19648189.2016.1217793.
URL http://dx.doi.org/10.1080/19648189.2016.1217793

[27] W. Van der Wiel, Probabilistic risk assessment of a system of dike ring areas, Master thesis, Delft University of Technology (2004).

[28] H. Verheij, Aanpassen van het bresgroeimodel in HIS-OM: Bureaustudie, Tech. rep., Deltares, Delft (2003).

[29] P. Kamrath, M. Disse, M. Hammer, J. Köngeter, Assessment of Discharge through a Dike Breach and Simulation of Flood Wave Propagation, Natural Hazards 38 (1) (2006) 63–78. doi:10.1007/s11069-005-8600-x.
URL http://dx.doi.org/10.1007/s11069-005-8600-x

[30] F. D. Foresee, M. T. Hagan, Gauss-Newton approximation to Bayesian learning, in: Neural Networks,1997., International Conference on, Vol. 3, 1997, pp. 1930–1935 vol.3. doi:10.1109/ICNN.1997.614194.

Appendix A. Failure probabilities of breach locations

This appendix provides additional information pertaining the modelling of the probability of failure of breach locations for the case as described in Section 2. First, we describe a (simplified) hydrodynamic model that describes the propagation of a (peak) discharge wave through the river branches. Then, once a breach occurs at a breach location, breach flows toward flood-prone areas are described as well as their interactions with the discharge flow through the river branches. Finally, we describe how this hydrodynamic model is used in conjunction with random variables in order to estimate failure probabilities of breach locations, and indicate how accurate these estimates are.

Appendix A.1. River model

The hydrodynamic model uses a discharge wave as input. This discharge wave is composed out of a normalized discharge wave as described in [4] and shown in Figure A.19. The discharge wave of Figure A.19 is scaled by multiplying the discharge wave with a sampled peak discharge (see Appendix A.3 for the sampling of peak discharges). These discharges are converted to water levels at each breach location using stage - discharge relations. The use of these stage - discharge relations is one of the key reasons why our model is computationally efficient; however, this computational efficiency comes at the cost of simplified river hydrodynamics.

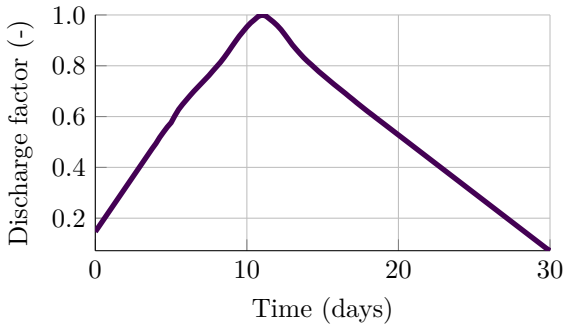

Figure A.19: Normalized discharge wave over time. The peak discharge, which coincides to a discharge factor of 1.0, occurs around day 11.

The stage - discharge relations are based on the Chézy formula using a single river profile as shown in Figure A.20. The Chézy formula is shown in Eq. A.1, which assumes equilibrium water levels:

$$Q_{riv} = 18 \log \left(\frac{12 \cdot A/C_{wet}}{k} \right) A \sqrt{A/C_{wet} \cdot i} \qquad \text{(A.1)}$$

where i is the slope, k is the Nikuradse coefficient, C_{wet} is the "wet" circumference of the river profile, A is the cross-sectional surface area and Q is the discharge. The river profile of Figure A.20 is used for the river Rhine. stage - discharge relations for the Pannerden Canal and IJssel are derived from the stage - discharge relationship for the Rhine. This derivation is based on two assumptions. The first assumption is that a constant fraction (1/3) of the discharge from the Rhine flows into the Pannerden Canal, and that a constant fraction (again 1/3) of the discharge from the Pannerden Canal flows into the IJssel. The second assumption is that given a discharge Q in the Rhine and a discharge $1/3Q$ in the Pannerden Canal, the water depth in both river branches will be identical.

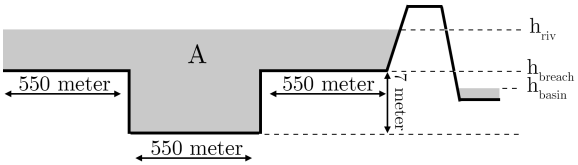

Figure A.20: Cross-sectional river profile as used for the river Rhine. The levels h_{riv}, h_{breach} and h_{basin} are relative to the bottom level of the river.

Based on the approach in [27], rather than iteratively solving Eq. A.1, a range of water levels and accompanying discharges is computed using Eq. A.1 and fitted (using a least squares method) to the following stage-discharge relation:

$$h_{riv} = a \cdot {Q_{riv}}^{b} \qquad \text{(A.2)}$$

Where a and b are found as a result of the fitting process. The resulting stage - discharge relations are shown in Figure A.21. These relations are based on the the river Rhine profile of Figure A.20, along with a Nikuradse coefficient k of 0.05 meter and a slope i of 10^{-4}.

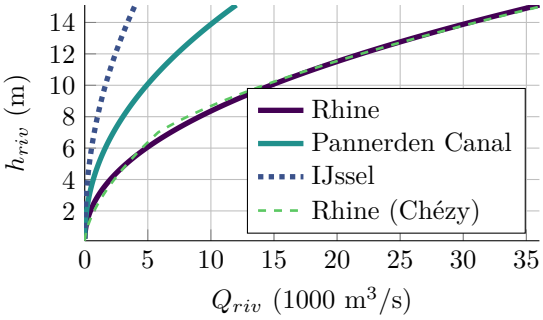

Figure A.21: Stage - discharge relationships for the three river branches using equation A.2. The Stage - discharge relationship for the river Rhine using the Chezy equation of Eq A.1 is shown as well.

Appendix A.2. Breach and flooding model

The stage - discharge relations of the previous section are used to determine whether or not a breach occurs at a breach location. If the water depth at a breach location exceeds the 'critical height' of that breach location (see also Appendix A.3), the breaching process will start.

The breach width B will then develop over time according to the Verheij - van der Knaap equation as shown in Eq A.3 [28], similar to (for example) [4, 5]. This equation is based on an analysis of a breached defence functioning as a (submerged) weir, connected to a basin where any inflow results in a constant increase of the water level over the basin surface (i.e. pumping mode assumption). We use a similar model for estimating the breach flow, and use an adjusted version of the Poleni weir equation (Eq. A.4) to estimate the breach flow Q_{breach}.

$$B = 1.3 \frac{g^{0.5} \left(h_{riv} - h_{basin}\right)^{1.5}}{u_c} \log\left(1 + \frac{0.04g}{u_c} t\right) \tag{A.3}$$

$$Q_{breach} = \frac{2}{3} m \sqrt{g} B \left(h_{riv} - h_{breach}\right) \sqrt{h_{riv} - h_{basin}} \tag{A.4}$$

In Eq. A.3, g is the gravitational acceleration (set to 9.81 m/s^2), t is the elapsed time since the initial moment of breach and u_c is a critical flow velocity. The critical flow velocity has been set to 0.3 m/s for the application in this study, which is close the values used for sand in [28]. Furthermore, the breach width B is limited to a maximum of 200 meters in accordance with [5].

Eq. A.4 is an adjusted version of the Poleni weir equation in which the addition of $\sqrt{h_{riv} - h_{basin}}$ accounts for submerged flow. Furthermore, m is a flow factor for energy losses and is assumed to be equal to 1. In a more realistic application, this factor could be estimated more accurately, for example using an approach as shown in [29].

The height of the water level in flood prone areas is estimated using a simple 'bath tub' model as shown in Eq. A.5.

$$\frac{dh_{basin}}{dt} = \frac{Q_{net}}{A_{basin}} \tag{A.5}$$

In Eq. A.5, Q_{net} is the net flow towards or out of the basin due to breach flows, with A_{basin} as the surface area of the basin. The surface area for the basin 'D48' is set to 400 km^2, while the surface areas for 'D49', 'D50' and 'D51' are set to 100 km^2.

The internal breaching between two flood prone areas (i.e. 'B49i' and 'B50i') is treated slightly different from breaches from the river, as these internal dikes are lower and generally weaker. It is assumed that, in case of an internal breach, the water will distribute immediately over both areas, in accordance with the pumping mode assumption of Eq. A.5.

Appendix A.3. Random variables

Random variables are introduced for both strength and load in the river model. The strength is represented by the critical heights of dikes at the breach locations, while the load is represented by the river discharge. Using the stage - discharge relations of Appendix A.1, the river discharge can be converted to a water level. If the local water level exceeds the critical height at a breach location, failure is induced and a breach will form.

The critical height (h_{crit}) or strength at a breach location is represented by a normal distribution. In the context of a reliability assessment, only the lower end of a strength distribution is of importance. This critical height represents the combined probability that the flood defence at a breach location fails due to various failure mechanisms (e.g. piping, overflow or macro-stability). The actual flood defence height (h_{def}) is assumed to be the mean of the normally distributed critical height, and will also be used as the input for determining the investment cost of a flood defence (see also Section 3.4).

The standard deviation of a critical height distribution is determined with a fixed coefficient of variation of 0.1. Furthermore, the internal breach locations are linked to their external counterparts by having a mean critical height at 90% of the external counterpart, similar to [4]. The mean of the critical height distribution of an external breach location for a desired failure probability is obtained by iteratively shifting the mean in Monte Carlo simulations (without hydrodynamic interactions) until it converges to the desired failure probability. This iterative shifting (regula falsi) of the mean is the same approach as described in [5].

The load is represented by a Gumbel distribution for the Rhine river discharge at Lobith (e.g. see [20]). The normalized discharge wave of Figure A.19 is multiplied with a peak discharge value. [20] investigated various importance sampling distributions. Based on that research, we use a uniform sampling distribution. The density functions of both the importance sampling and regular distribution are shown in Figure A.23.

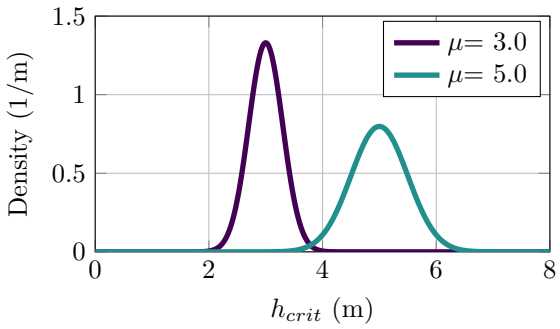

Figure A.22: Density functions of two normally distributed critical heights with a coefficient of variation of 0.1 and means of three and five meters, respectively.

Appendix A.4. Variability in failure probability estimates

For reproducibility, the randomness of the failure probability estimation itself was removed by fixing the seed of the random number generation before each Monte Carlo simulation.

In order to obtain insight in the variability of failure probability estimates without resetting the seed, the fail-

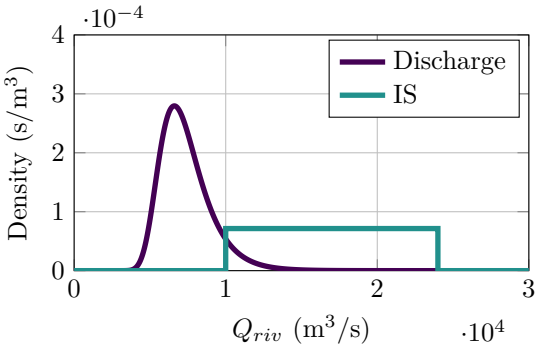

Figure A.23: Density functions of the Gumbel distributed discharge of the river Rhine, and the uniformly distributed importance sampling (IS) function.

ure probability was estimated for a range of sample sizes. The sample sizes tested range from 100 to 10,000 samples, which each were repeated 10,000 times to estimate the variability of each sample size. A sufficiently small variability is defined here as a maximum of around 10% over- or underestimation of the average estimate of the failure probability. It can be expected that as the number of samples goes up, the variability of the failure probability estimate goes down.

The failure probability results from evaluation of the simplified limit state function in Eq. A.6.

$$Z = h_{def} - h_{riv} \tag{A.6}$$

Where the load is a peak water level in a river (h_{riv}, in meters) and the strength is represented by critical height of the flood defence (h_{def}). h_{riv} is converted from a Gumbel distributed discharge (sampled using the Importance Sampling strategy described in Appendix A.3). The conversion from discharge to peak water level (h_{riv}) uses the conversion described in Appendix A.1. The strength is taken from a normal distribution with a CoV (coefficient of variation) of 0.1 and with a mean of three meters or five meters. The means of three and five meters represent the upper and lower bound of the range of expected strength in the economic optimisation; see also 3.5.

The results are shown in Figure A.24. The range which includes 95% of the failure estimates converges quicker for a mean of three meters than a mean of five meters in Figure A.24. Based on these figures, 5000 samples was seen as a sufficient number of samples, as the 95% variability range falls either within the ± 10% threshold, or closely approaches this threshold.

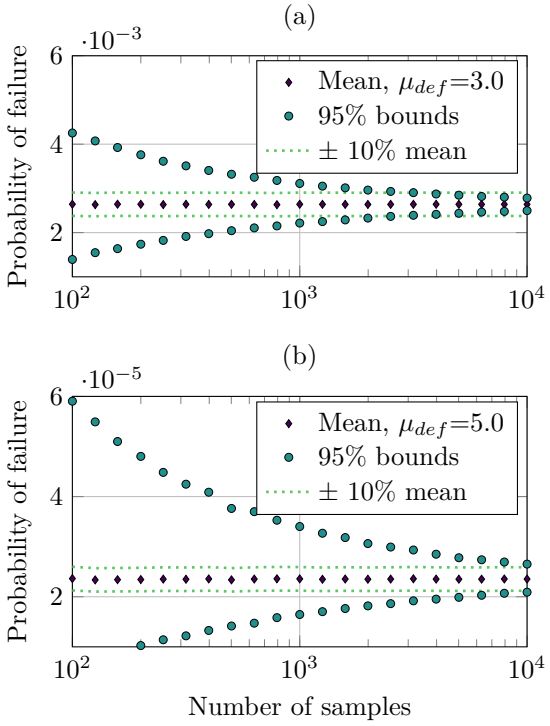

Figure A.24: Example of the variability of a failure probability estimate (using Eq. A.6) as a function of the number of samples. The load has a Gumbel distribution, while the strength has a Normal distribution with a mean (μ_{def}) of three meters (a) or five meters (b).

Appendix B. 'Goodness of fit' of the surrogate model

A neural network was trained for each flood prone area in the case study area. For each neural network, the number of neurons on the input layer are the correlated breach locations as mentioned in Table 3. The output layer contains a single neuron which represents the EAD in that area. The neural network structures as described in Table B.4 were used.

Table B.4 shows the fit (R^2) and the number of neurons in each hidden layer. The number of layers and/or number of neurons needs to be increased for the downstream areas in order to attain a similar R^2, which indicates that the response of the flood risk cost is more complex for these areas: the further downstream a flood prone area is, the more flooding scenarios are possible.

With surrogate modelling, there is a possibility of overfitting. Overfitting is defined as having a very good fit on the training data, but as soon as new data is presented, poor fits are achieved instead. A few techniques were used to prevent this. First, the training was done using 70% of the data, with 15% used to test network generalization and another 15% as an independent measure of network performance. Secondly, Bayesian regularization was used to improve network generalization [30]. Nevertheless, even with the very high R^2 values, the neural nets do not always show accurate, expected behavior; compare for example the two lines in Figure B.25.

In order to test the accuracy of the neural networks, the correlation coefficients of Table 3 were re-calculated using the output of the fitted neural networks with 823,543 system configurations (7^7) of critical dike heights. The resulting correlation coefficients are shown in Table B.5. Comparing Table 3 & B.5 shows that there are only minor differences between the calculated correlation coefficients. The largest difference is found for the correlation coefficient between breach location B482 and flood prone area D49, which is 0.09 in Table 3 and -0.02 in Table B.5. Nevertheless, both are weak correlations.

Area	Neurons in HL1	Neurons in HL2	R^2
D48	30	0	0.99999
D49	40	0	0.99991
D50	30	3	0.99994
D51	30	4	0.99990

Table B.4: Amount of neurons in each hidden layer (called HL1 and HL2) in the neural nets used for approximating the modeled risk cost in D48, D49, D50 and D51. The R^2 is indicative of how well the neural net approximates the modeled data.

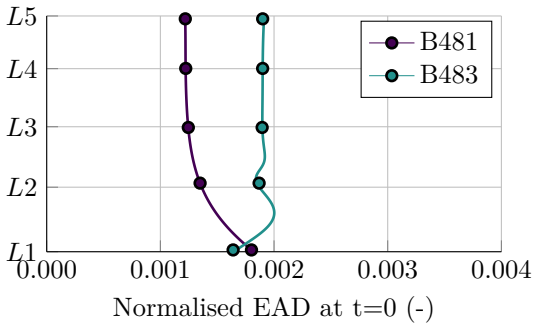

Figure B.25: Examples of expected and unexpected behavior of the neural network; the 'wobbling' between L1 and L3 for line B483 is unexpected. Shown is the normalised EAD of D48 versus the level of breach location B481. Also shown is the normalised EAD of D49 versus the level of breach location B483. Both lines assume that all other breach locations are kept at their L1 level and are, and are normalised on the maximum flood damage of the associated flood prone areas.

Breach location	D48	D49	D50	D51
GER	0.01	-0.54	-0.50	-0.37
B481	-0.67	0.06	0.05	0.05
B482	-0.48	-0.02	0.09	0.08
B483	-0.29	0.12	0.14	0.11
B49	0.00	-0.75	0.09	0.17
B50	0.00	0.00	-0.78	0.01
B51	0.00	0.00	0.00	-0.84

Table B.5: Spearman's correlation coefficients (calculated with the surrogate model) for the critical heights of breach locations (rows), versus the annual flood risk cost of the flood prone areas (columns).