Bayesian networks for levee system reliability

Reliability updating and model verification

Roscoe, Kathryn

DOI
10.4233/uuid:9968b155-539f-4e40-9562-5996a2843aa8

Publication date
2017

Document Version
Final published version

Citation (APA)
Roscoe, K. (2017). *Bayesian networks for levee system reliability: Reliability updating and model verification*. [Dissertation (TU Delft), Delft University of Technology]. https://doi.org/10.4233/uuid:9968b155-539f-4e40-9562-5996a2843aa8

BAYESIAN NETWORKS FOR LEVEE SYSTEM RELIABILITY

RELIABILITY UPDATING AND MODEL VERIFICATION

BAYESIAN NETWORKS FOR LEVEE SYSTEM RELIABILITY

RELIABILITY UPDATING AND MODEL VERIFICATION

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op 17 februari 2017 om 10:00 uur

door

Kathryn Lynn ROSCOE

Master of Science, Hydrologic Sciences, University of California, Davis
geboren te Miami, Florida, Verenigde Staten van Amerika.

This dissertation has been approved by the promotors:

Prof. dr. ir. J.K. Vrijling
Prof. ir. A.C.W.M. Vrouwenvelder

Composition of the doctoral committee:

Rector Magnificus	Chairman
Prof. dr. ir. J.K. Vrijling	Delft University of Technology, promotor
Prof. ir. A.C.W.M. Vrouwenvelder	Delft University of Technology, promotor
Independent members:	
Prof. dr. R.M. Cooke	Resources for the Future, USA
Prof. dr. ir. P.H.A.J.M. van Gelder	Delft University of Technology
Prof. dr. D. Straub	Technical University Munich, Germany
Dr. ir. T. Schweckendiek	Delft University of Technology and Deltares
Prof. dr. ir. S.N. Jonkman	Delft University of Technology, reserve member
Other member:	
Dr. A.M. Hanea	University of Melbourne, Australia, supervisor

Keywords: Levee reliability, system reliability, failure probability, Bayesian networks, proven strength, survival observations, correlated components, length effect, Bayesian updating, model uncertainty

Printed by: Druk. Tan Heck, Delft

Front & Back: Kathryn Roscoe

Cover photo: Henri Cormont, Beeldbank Rijkswaterstaat

ISBN 978-90-6824-059-7

An electronic version of this dissertation is available at
http://repository.tudelft.nl/.

Probability theory is nothing but common sense reduced to calculation.

Pierre Simon Laplace

CONTENTS

SUMMARY

Flood risk analysis is necessary to make smart, informed decisions about which risk reduction measures deserve priority. When levee systems play a key role in flood protection, these decisions often translate to which levee improvements should be carried out first. In flood risk analysis, the probability that a levee system fails is a critical component, but one that is wrought with uncertainty. Much research has focused on how to calculate the probability of system failure. However, for levees, what is typically seen in practice is a simplification of the system to make calculating the system failure probability easier.

In the Netherlands, over 30 years of research has led to a rigorous methodology for calculating the probability of levee system failure, which has been encoded into the software Hydra-Ring. Two key algorithms calculate (1) the segment failure probability and (2) the system failure probability. The first is referred to in this dissertation as the modified outcrossing (MO) method, and takes into account the spatial autocorrelations within a levee segment. The latter, referred to as the Equivalent Planes (EP) method, accounts for the correlation between levee segments. The methods are both approximate, and very efficient, but a thorough description of them, as well as a verification, was lacking in the literature. Furthermore, there has been a surge of interest recently in using survival observations - the survival of a levee during an observed (high) water level - to update levee reliability estimates. However, use of the MO and EP algorithms in combination with updating has not been explored. The implementation and accuracy of these algorithms in combination with a survival observation are topics of current relevance.

This dissertation explores the development and use of a Bayesian network (BN) for levee system reliability, to augment and verify the methods already in use in the Netherlands. BNs are a type of probabilistic graphical model, in which correlations between variables can be seen in the structure of network. The BN selected for use in this dissertation works with Monte-Carlo (MC) sampling, and correlates variables in the network using the Gaussian copula. In this sense, it can be considered a more explicit, less approximating method than the algorithms in Hydra-Ring. The BN was used to test the MO algorithm, and MC directional sampling and exact solutions were used to test the EP algorithm. While both methods produce some error relative to more exact MC methods, the error is not substantial, even after incorporating a survival observation. The BN was applied to two case studies in the Netherlands, to calculate system failure probabilities due to the piping failure mechanism. In these cases survival observations were used to improve the system reliability estimate. These applications show that not all survival observations have equal impact on the levee system reliability estimate. It was investigated under which conditions survival observations are useful. A BN was also developed specifically for the estimation of the model uncertainty in a geotechnical failure model. This uncertainty can dominate the failure probability estimate, and it is therefore important to estimate it as sharply as possible. The research in this dissertation shows that us-

ing a BN, high quality hindcasts (geotechnical model output for historic input data) can be used together with observed failure (or survival) to substantially improve the model uncertainty estimate, even with limited data.

The BN developed in this dissertation serves as a useful augmentation to the levee system reliability methods currently in use. Its computation time is not prohibitively slow, but it can be hundreds of times slower than the approximate algorithms within Hydra-Ring. Therefore, it should not be seen as a replacement for Hydra-Ring, but rather a yardstick which can be used to verify Hydra-Ring algorithms when results are questionable, or when survival observations are expected to be useful.

ACKNOWLEDGMENTS

The experience of this PhD makes me feel like I can relate to regular people who go out and decide to become competitive weight-lifters. The sheer will that I had to find to overcome rough patches (both technical and emotional) makes me want to melt into a puddle of exhaustion just remembering it. But I didn't do it alone. I have such immense gratitude towards a number of people, who served as rods of support without whom this PhD never would have come to fruition.

First, I want to acknowledge my professors, Han Vrijling and Ton Vrouwenvelder. You are giants in the field of flood risk and reliability modeling, and it has been a privilege to see the reliability world through your eyes, both the technical specifics, as well as the broader context. One would be hard-pressed to find two men more dedicated and invested in their field, working on weekends and at night despite having retired. It's inspiring to see, and I am so grateful for your time, your guidance and perspective.

My supervisor, Anca Hanea, has stood by me, challenged me, encouraged me, improved me, for absolutely no personal gain. She didn't know me when I approached her to be my supervisor, but she took me on and her support and commitment never wavered, even when she moved to the other side of the world. She is a model to me of selflessness, strength, authenticity, and empathy. I am forever and eternally grateful for your support.

I entered the world of risk and reliability through my work at Deltares, which began in 2008. Throughout my PhD they offered support, in all forms, at every phase. For that I am deeply grateful. I want to thank Gerard Blom, Jaap Kwadijk and Annemargreet de Leeuw, who helped me bridge the roughest patch, and whose belief in me had a substantial impact on my decision to see this thing through. I would also like to express my gratitude to Wim Kanning, who - while inundated in work - made time to meet with me almost every Friday to discuss and advise on my research as well as giving me countless tips on practical aspects of the PhD. At the TU Delft, I want to thank Gerrit Schoups for numerous enlightening conversations about Bayesian networks, and particularly for helping shape my ideas for the model error research. Your genuine enthusiasm for research is inspiring. I would like to thank the TU Delft department of Hydraulic Engineering and Bas Jonkman, for their financial and logistical support. I am also grateful to STW - the funding agency who made my research possible. And to Matthijs Kok, who headed the program my PhD was a part of, I am so thankful. He brought my research ideas, which I had been eager to explore for years, into his program, and without that help, these ideas would have probably remained in my head. He worked tirelessly to manage the research program, and always with kindness and humility.

And now, a bit more on the personal side. To my mom, my number one fan. Nobody in this world thinks I'm smarter or more beautiful than my mom. She has been my cheerleader throughout this PhD, encouraging me, and pushing me to keep on it. She has flown here countless times, staying for weeks at a time to watch Lucas and lighten

our load. She would do anything for me, she makes my problems her problems, and she is always on my side. I can never thank her enough for that.

And to friends, colleagues, and family who lifted me up, distracted me, let me vent, made me laugh, helped me in my ongoing journey to figure out how and who I want to be. To my sister, for helping me with good advice at the lowest point in this PhD. To my mom-group, particuarly Karen, Elizabeth, Caroline, Laura, Dani, and Holly, for making Delft really feel like my home, for eliciting cathartic bouts of laughter, for giving me unending craft ideas. To my colleagues at Deltares - Marjolein, Janneke, Kristin, Nienke - for coffee breaks and lunch walks in which venting about just about anything was allowed. To Lauren, my soulmate friend whose existence in this world is like sunshine on closed eyelids; I savor your sharp wit, humor, and perspectives on everything. To all of you: thank you for just being you and making my life fuller.

To Karen and Wendy, without whom I feel absolutely sure I never would have finished this PhD. They took care of Lucas when I couldn't, in such a loving, playful, educative way. I wanted to take care of him myself, and my heart nearly broke from having to program Matlab code instead of making leaf crafts with Lucas. But the pictures they'd send of him having so much fun with them made it possible to focus on work, in a way that a day care never could. Thank you both from the depths of my soul, for loving and caring for my favorite little person in this world.

My pillar of support, without whom I never could have done this, is Ferdinand. He was all things to me in this process: mentor, sounding board, reviewer. And at every point when I thought about quitting, he never judged. He just helped me decide for myself what was the best path for me, and whatever I decided he would throw his whole weight behind. For his unconditional love and support, for the much needed laughter his humor elicits in me, for the 3 million dinners he's made and the 4000 times he's cleaned the kitchen, for absorbing the wild oscillations in my mood/perspective on life/thoughts about everything with humor and stability, for all of this and so very much more, I thank you with all of my heart.

And last, my kid. My heart and soul, Lucas. I have learned more about myself and human nature these past four years than anything else. It is a great adventure to have a child. An exercise in exhaustion and refining patience, yes. But also a fulfillment that eclipses anything I'd ever known before. A joy so encompassing when he catapults into our bed in the morning. His existence has helped me put this PhD in perspective. It is an accomplishment of which I am so proud. But it pales in comparison to how proud I am of our absolutely irrational, feverishly curious, comically stubborn peanut butter-spinach-pumpkin-pla-pla-Lukie Pukie-munchkin pie. This PhD is really dedicated to him.

1

INTRODUCTION

1.1. FLOOD RISK AND LEVEE SYSTEM RELIABILITY

Floods are the most common natural disaster worldwide. According to the 2015 World Disasters report, floods represent 45% of all natural disasters over the period 2005-2014. They are dangerous, potentially lethal, and extremely costly, not to mention the psychological damage they inflict on victims. Each year (on average) during the period 2005-2014, 5,900 people died due to floods, 87 million people were affected, and 34 billion U.S. dollars of damage was incurred.

Levee systems are a common preventative strategy to reduce flood risk. Often, several sections of a levee system may be in need of improvement, but limited financial resources make it impossible to improve them all. In the past two decades there has been a strong move towards risk-based approaches of flood management, in which both the probability of flooding and the consequences are accounted for. Risk analyses allow decision-makers to best determine where to allocate funds. For example, a weak levee that is protecting agricultural land will probably have less risk associated with it than a levee in moderate condition that protects residential communities or a business district. Knowing the risk associated with the reaches in a levee system allows flood managers to prioritize improvement measures. However, risk estimates are usually subject to large uncertainty ([1], [2], [3]), which makes it difficult to come to policy decisions or decide on improvement measures ([4]). One of these sources of uncertainty is the reliability of levee systems, which is the focus of this dissertation. A poor estimate of the reliability of a levee system leads to a poor estimation of the flooding probability, and in turn, a poor estimate of risk.

The reliability of a levee is the probability that it will maintain its flood protection function; it is the probability that the levee will *not fail*. It is related to the failure probability (P_f) of the levee in that it equals $1 - P_f$. A levee system is a long spatial extent of levees, usually comprising a number of statistically homogeneous segments. When calculating the risk of a certain protected area, the probability of failure of the entire system of levees protecting that area must be calculated. Figure 1.1 presents an illustrative fault

1

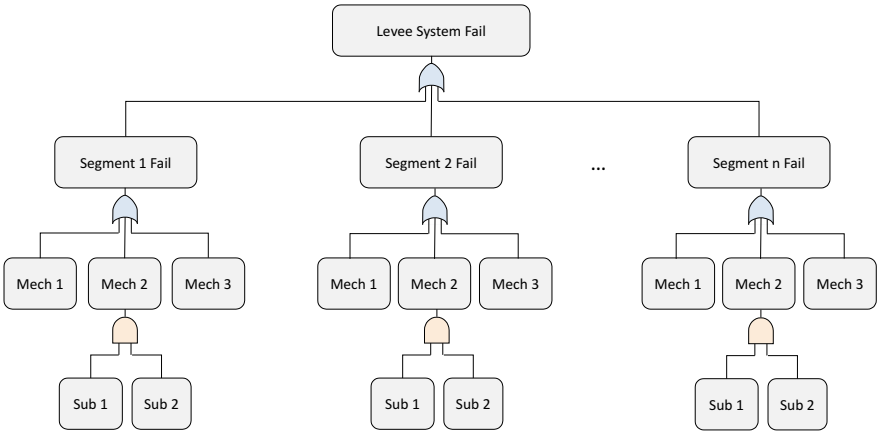

Figure 1.1: Example of a fault tree, showing the components of levee system reliability, and the 'and' (orange) and 'or' (blue) gates. In this illustration, the system fails if any of the segments fail. Each levee segment can fail if any of three mechanisms occurs (Mech 1, Mech 2 and Mech 3). The second mechanism is a system of two sub-mechanisms (Sub 1 and Sub 2) connected by an 'and' gate, which means both must occur in order for Mech 2 to occur.

tree diagram, which shows the connection between components of the system and the system failure probability. In the illustration, system failure is connected to the segment failures via an OR gate. This means that if *any* of the segments fails, the system will fail. Similarly, each segment failure is connected to three failure mechanisms, also by an OR gate, meaning that if *any* of the mechanisms occurs, the segment will fail. One of the failure mechanisms in the illustration is further connected to two sub-mechanisms via an AND gate. This means that *both* sub-mechanisms must occur for the failure mechanism to occur. Section 1.2 describes methods to calculate the probability of such systems. Some basic concepts in reliability analysis that are relevant to this thesis are provided in Appendix B.

1.2. METHODS TO CALCULATE LEVEE SYSTEM RELIABILITY

System reliability is notoriously complex to compute, due to the interdependence of system components. In levee system reliability, the interdependency arises from the shared soil and load variables (for dependence between failure modes), and their spatial correlation (for dependence between segments). The system failure probability is only straight-forward to calculate in the trivial cases in which all of the components are either fully correlated or entirely independent - which is almost never the case in levee systems. Appendix A provides background about why calculation of the system failure probability becomes more challenging when components are correlated.

In the last two decades, a number of methods have been developed to calculate system reliability when the system components are correlated. A detailed overview of these methods is given in Chapter 4. However, such methods are rarely used in levee reliability. Rather, what is often done in practice is some simplification of the system into the trivial cases of fully dependent or completely independent components. For example, in

an extensive methodology to compute flood risk at a national scale in the UK ([5]), three assumptions were made to avoid the complexity of spatial variability: (1) the loads are fully correlated over all levee sections, (2) the resistance is independent between levee sections, and (3) the resistance is fully correlated within a levee section. The first two assumptions are reasonable, but the third is dubious, and can lead to underestimating the failure probability. An aspect of levee system reliability that many practitioners grapple with is the length effect, which looks at how the failure probability of a segment increases relative to a cross-section. The larger failure probability at the segment level is the effect of spatial variability within a statistically-homogeneous segment of levee, and the likelihood of finding a 'weak spot'. This is described in detail in Chapter 3. When the length effect is accounted for (which in many cases, such as in the UK study cited above, it is not), analysts tend to simplify the problem, again to avoid the issue of correlated components. Examples of this are given in the introduction of Chapter 3 .

The most rigorous methodology for computing levee system reliability was developed as part of a national flood risk assessment in the Netherlands, described in [6]. The Netherlands has a long history in probabilistic design and assessment of flood defense systems, starting in the 1980s ([7]). A history of probabilistic flood defense reliability modeling in the Netherlands, which really has its roots in the catastrophic flood of 1953, is provided in Appendix C. In the late 1990s, the techniques that had been developed were encoded in a software package known as PC-Ring, which calculates the reliability of large flood defense systems (which may include levees, dunes, and hydraulic structures), subjected to various loads (lakes, sea, and rivers); it also considers multiple failure mechanisms. Starting in 2010, the methods in PC-Ring were recoded in the reliability software Hydra-Ring, which has a more flexible format and includes additional functionality. Overviews of PC-Ring can be found in [8] and [9]. More details about how temporal and spatial variability are accounted for in the model can be found in [10]. For a very thorough description of the methods and algorithms within Hydra-Ring, the reader is referred to the Hydra-Ring technical document ([11]). Because some of the algorithms in the model have a central role in this dissertation, a brief overview of the model is provided in the following section.

FLOOD DEFENSE RELIABILITY MODEL PC-RING/HYDRA-RING

The PC-Ring/Hydra-Ring model calculates failure probabilities for levee systems in a piece-wise fashion. It first computes component failure probabilities, and then engages a number of combining or upscaling algorithms which combine the component failure probabilities to arrive at the system failure probability. The model accounts for a lot of complexity, including:

- Loads that vary at different temporal scales (e.g. water level at sea vs. discharge in a river)
- Influence of storm surge barriers
- Statistics that vary per wind direction (e.g. wave heights at sea)
- Spatial variability within and between statistically homogeneous levee reaches

1

An initial component in the PC-Ring/Hydra-Ring model is defined as a single cross-section within a levee segment, a single wind direction, a small time scale (the time scale of the fastest fluctuating load variable), a single failure mechanism, and a single state of any closure barriers. These components are then iteratively combined. The order in which they are combined depends on the way that the temporal variability of the load is modeled. For a 'block model', also referred to as an FBC model for Ferry Borges-Castanheta ([12]), the sequential steps in the algorithm are listed below, and follow loosely from the Hydra-Ring Technical Reference Manual (with some extraneous detail removed). The steps which are relevant in this dissertation (see boxed text in list below) are the upscaling of the probability from cross section to levee segment (step 3), and the combining of the failure probabilities of the failure mechanisms and segments (step 6). The former is calculating the length effect; the algorithm is referred to in this dissertation as the Modified Outcrossing method. The latter is the algorithm to combine correlated components, referred to as the Equivalent Planes method. Very little about both of these algorithms has been published or verified, which is one of the focal points in this dissertation.

1. Determine the failure probability of the smallest component: a single cross section, a single failure mechanism, a single wind direction, one closure situation and a small time increment.
2. Combine failure probabilities of the closure situations.
3. Upscale the failure probability of a cross section to a levee segment
4. Combine the failure probabilities over the wind directions.
5. Upscale the failure probabilities temporally to a year.
6. Combine failure probabilities of all failure mechanisms and segments.

There are a number of references which describe the PC-Ring model. In 2003, Lassing, Vrouwenvelder, and Waarts published a paper in which the general overview of the model is provided, but where the focus is primarily on the failure mechanisms that are included in PC-Ring ([9]). A year later, Steenbergen, Lassing, Vrouwenvelder, and Waarts published a complimentary paper in which the focus was on the probabilistic methods in PC-Ring ([8]). Both papers include an example application. Vrouwenvelder published a paper in 2006 in which some of the concepts in the previous papers are reiterated, but where there is more focus on how PC-Ring handles spatial and temporal variability. In 2015, Jongejan and Maaskant published a paper about risk analysis in the Netherlands ([6]); in it, the writers clarify how PC-Ring fits into the bigger picture of calculating risk. These papers help give glimpses into the model, its complexity, its applicability, and some of the algorithms. However, details tend to be lacking, and a definitive resource on PC-Ring was only made available in a series of Dutch reports. Hydra-Ring, by contrast, has been carefully documented in an English-language technical reference manual. It is a thorough resource for all the programming specifics, the probabilistic algorithms, hydraulic models, and failure models contained within Hydra-Ring. For details on any

of the aspects of Hydra-Ring not found in this dissertation, the reader is referred to this technical reference manual ([11]).

1.3. PROBLEM STATEMENT AND RESEARCH OBJECTIVES

Flooding is the most relevant natural hazard threat in the Netherlands. The large fraction of the country that is flood prone, coupled with the economic activity that occurs in those flood-prone regions, makes the consequences of flooding severe. The state of the flood defenses is quite literally a matter of national security. To this end, the government has mandated that all primary flood defenses be assessed on a recurring basis to ensure they meet strict protection standards. The research and development of tools to support the assessments receive millions in government spending, and the tools are in a continuous state of revision and improvement. The reliability model Hydra-Ring, described in Section 1.2, is at the heart of the new set of tools for the national flood defense assessments in the Netherlands. However, some of the innovative algorithms in the model, which make it so efficient, have never been tested for accuracy.

The overarching problem that is addressed in this PhD is the lack of confidence that can arise in computed failure probability estimates of levee systems. Generally, this happens when the intuition or experience of those familiar with the system is at odds with the calculated failure probability. Intuition about the strength of a levee is fed by observations of either good or poor performance under higher-than-average loads. Because safety standards in the Netherlands are expressed as very low acceptable failure probabilities (e.g. annual probability of 10^{-4}), observations in the duration of record (usually no more than 150 years) are insufficient to prove that the levee satisfies this standard. However, they are sufficient to build intuition. For example, when the failure probability estimate is high, but there has been no evidence of weakness of the levee, managers and others may become skeptical of the estimate. Similarly, if sand boils, cracks, or other signs of weakness have been observed, a very low estimated failure probability would be suspect. Poor estimates of the failure probability can be caused by a number of sources. Uncertainty in the soil parameters can lead to poor (or overly wide) prior distributions. The output of the geotechnical failure models which use these parameters are subject to uncertainty that is difficult to quantify. When failure probability estimates are produced which contradict observations, there is also a sense of concern about the unverified algorithms in Hydra-Ring.

The research in this dissertation focuses on the development of a Bayesian network (BN) for levee system reliability that can use evidence to improve reliability estimates and reduce uncertainty in the soil parameters, at the system scale. It also focuses on the verification of two of the key algorithms in Hydra-Ring, one of which combines correlated components in a system, and the other which accounts for the length effect. The BN developed in this research plays a key role in the verification of this latter algorithm. Another issue addressed in the dissertation is the uncertainty in geotechnical failure models, which are used to determine if, for a given failure mode like piping or macrostability, a levee can resist a particular load. These models are an integral part of the reliability analysis. In the national flood defense assessments carried out in the Netherlands, the model error is currently based on expert opinion; observations are not explicitly taken into account. The research in this dissertation considers the use of a

BN, together with geotechnical model hindcasts and levee performance observations, to quantitatively estimate the uncertainty in failure model output.

To summarize, the questions the research presented in this dissertation strives to answer are enumerated below. The chapter(s) in which the question is addressed is provided in brackets.

1. Is it possible to develop a BN that can handle the requirements of levee system reliability? [Chapter 2]

2. How can the BN be applied to update the failure probability of a levee system, given survival observations? [Chapters 2 and 5]

3. How does the efficiency of the BN compare with the efficiency of the algorithms in Hydra-Ring? [Chapters 3 and 5]

4. How accurate is the Hydra-Ring algorithm which combines correlated components? [Chapters 4 and 5]

5. How accurate is the Hydra-Ring algorithm that accounts for the length effect? [Chapters 3 and 5]

6. Are Hydra-Ring algorithms able to be used to perform inference, and are the results accurate? [Chapters 3 and 5]

7. Can a BN be used, together with geotechnical failure model hindcasts and levee performance observations, to estimate the error in a failure model? [Chapter 6]

1.4. RESEARCH APPROACH

The first part of the PhD research is the development of a general methodology to apply a BN to levee system reliability. BNs are becoming increasingly popular for modeling uncertain systems with high complexity, but have not yet been applied to the problem of levee system reliability. They are a form of graphical model, which means the variables in the system and the dependence between them is given visually. This helps simplify the joint distribution, because the dependence is encoded in the graphical structure, and the variables are then specified by marginal or conditional probability distributions. A well-known feature of BNs is their ability to perform inference, which means that information about any part of the system can be used to improve our understanding of the rest of the system. Such a capability is highly relevant in levee reliability modeling, where we are burdened by a great deal of uncertainty, but often observe the performance of the levee (e.g. for a given load, we see that the levee performed well (no damage), or poorly (cracks, sand boils, etc)). A BN allows us to use this information to reduce our uncertainty in system variables (like soil cohesion, friction, porosity, etc), and ultimately to improve the reliability estimate.

Following development of the BN methodology, the BN is used to verify the length effect algorithm in Hydra-Ring (referred to in the dissertation as the modified outcrossing method) via numerical examples. The algorithm which combines correlated components - referred to herein as the Equivalent planes method - is verified via a wide range of

numerical (synthetic) examples, using exact solutions or Monte-Carlo directional sampling as a reference calculation. The use of the BN in this context is also explored and compared with the Hydra-Ring algorithm. The BN methodology is further applied to two case studies in the Netherlands to calculate and update the system failure probability due to the piping mechanism. In the case studies, survival observations (i.e. observed water level and no evidence of failure) are used to improve the system reliability estimate as well as develop the posterior joint distribution of the soil parameters. In one of the cases, the Hydra-Ring algorithms are applied in addition to the BN, to test the accuracy of the system failure probability calculated with the Hydra-Ring algorithms in a real application, both prior to and following incorporation of a survival observation. The last issue addressed in the thesis is the challenge of estimating the uncertainty distribution of failure model output (e.g. the stability factor estimated by a slope stability model). A BN is developed specifically for this case, and hindcasted model results are used together with failure and survival observations to estimate posterior uncertainty distributions. The method is applied to synthetic data to explore the goodness of the methodology under different conditions, and is also applied to a case study using macrostability hindcasts and observed performance of levees in the Netherlands.

MC-BASED BAYESIAN NETWORK

Modeling the reliability of a levee system with a BN is an alternative method to the methods in Hydra-Ring. Part of the research explores the use of both methods to address the same problem, in part to verify the Hydra-Ring methods (which contain more approximations), and in part to compare efficiency of the methods. As will be described in Chapter 2, the BN used in this dissertation works with Monte Carlo (MC) sampling, and applies the correlation structure defined in the network using the Gaussian copula. For the type of applications in this dissertation, the method could be considered a copula-based MC approach. However, there are a number of advantages to the BN framework, which are addressed in Chapter 2, Section 2.4. Still, the comparison between the methods in Hydra-Ring and the BN method can be considered a comparison between an approximative method on the one hand, and an explicitly-modeled MC method on the other.

1.5. STRUCTURE OF THE DISSERTATION

This dissertation is structured as follows.

Chapter 2 presents the BN methodology developed as part of the current research for calculating and updating the reliability of levee systems. In this chapter, background about BNs is presented, their use in reliability estimation, and details about the specific type of BN that is used in this dissertation.

Chapter 3 focuses on the spatial aspects of levee system reliability, and uses the developed BN to test the goodness of a key algorithm in the levee system reliability models PC-Ring and Hydra-Ring: the modified-outcrossing method which computes the length effect. The accuracy is investigated for both prior and posterior estimates of the reliability, the latter resulting from survival observations; that is, a high load which a levee survives.

Chapter 4 investigates the accuracy of another key component in the PC-Ring and Hydra-Ring models: the Equivalent Planes method, which is used to calculate the system reliability of correlated components. In this chapter, exact solutions and MC approaches are used to test the goodness of the algorithm. The alternative use of the BN is also explored.

Chapter 5 presents two applications of the BN, to a regional and primary Dutch levee system. The applications focus on the piping failure mechanism, with survival observations. In each case, specifics of the application highlight implementation issues and solutions for practical use. In the case with primary levees, results are compared with those derived using the algorithms in Hydra-Ring and PC-Ring.

Chapter 6 develops a BN methodology that is specific for estimating the uncertainty in geotechnical failure models, based on observations. The uncertainty in these models can often overwhelm a reliability analysis, essentially dominating the influence on the failure probability. Historically, the uncertainty in the model output has been estimated based on the outcome of expert meetings. This chapter presents a quantitative method to assess and reduce the uncertainty based on observations.

Chapter 7 provides summarizing and concluding remarks, as well as suggestions for further research.

2

GENERAL METHODOLOGY: BAYESIAN NETWORKS IN LEVEE SYSTEM RELIABILITY MODELING[1]

2.1. INTRODUCTION TO BAYESIAN NETWORKS

Bayesian networks are an intuitive way to model multivariate probability distributions. They are a form of graphical model called directed acyclic graphs (DAGs), and consist of nodes, which represent random variables, and arrows, which indicate dependence between variables. The idea of a Bayesian network is to simplify a multivariate integral by coding the dependence via the graphical structure, and letting each variable be represented by a (conditional) probability distribution. For the example depicted in Figure 2.1, the Bayesian network consists of three random variables: X_1, X_2, and X_3. In this case, X_3 depends on X_1 and X_2; X_1 and X_2 are referred to as the parents of X_3 (denoted $pa(X_3)$), and X_3 is referred to as the child of X_1 and X_2. Because X_1 and X_2 have no parents, they are called root nodes. The Bayesian network simplifies the joint distribution representation by making use of the independence information in the graph (see Equation 2.1)

$$P(X_1, X_2, X_3) = P(X_1) \cdot P(X_2) \cdot P(X_3|X_1, X_2) \tag{2.1}$$

More generally, the joint density of any n variables is:

$$f_{1...n}(x_1, ..., x_n) = \prod_{i=1}^{n} f_{i|pa(i)}\left(x_i|pa(x_i)\right) \tag{2.2}$$

where $f_{1...n}$ is the joint density over the variables x_1 to x_n, and $f_{i|pa(i)}$ is the conditional probability of x_i given its parents.

[1] Parts of this chapter have been published as: Roscoe, K., and Hanea, A. (2015, July). Bayesian networks in levee reliability. In *ICASP12: 12th International Conference on Applications of Statistics and Probability in Civil Engineering, Vancouver, Canada, 12-15 July 2015.*

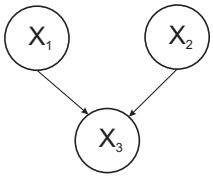

Figure 2.1: Three-variable Bayesian Network; X_1 and X_2 are the parent nodes and X_3 is the child node

What makes Bayesian networks particularly attractive is that they can be used to perform inference, which means if a variable anywhere in the network is observed, we can compute updated probability distributions of all the other variables in the network ([13]). This is very useful when observations are available and there is a lot of uncertainty in the variables in the network. Further, the intuitive graphical structure of a BN makes it easy for interested parties without a background in probability theory to understand the model, while maintaining a rigorous mathematical basis ([14]). Another advantage of the graphical structure, is that the dependence or independence of variables in the network can be read from the structure of the network, using a concept known as D-separation. This is important because once one or more variables in the network are observed, previously independent variables can become dependent, or vice versa. When using posterior distributions obtained from inference, it is important to take posterior dependence into account. The concept of D-separation is described in more detail in Appendix D.

Efficient exact and approximate algorithms are available for computing inference in Bayesian networks ([15], [16], [17]) but many rely on the nodes being represented by discrete (conditional) probability tables. For many applications this is not a limitation (consider genetics where variables like chromosomes and blood type have discrete states). However, when variables have a continuous distribution (as they typically do in engineering reliability applications), they must be discretized to make use of traditional inference algorithms. Research has focused on clever methods to perform discretization ([18], [19], [20]), but ideally it would be avoided altogether. This is because in reliability analysis, the tails of the distribution are so important - the low-probability high-impact values. To properly approximate the tails via discretization causes the number of discretization bins to become too large, and the computational effort too burdensome, to be feasible ([21]). An alternative to discretization is the Mixture of Truncated Exponentials ([22], [23], [24], [25]), which breaks any continuous distribution up into truncated exponential distributions. This method shows promise at being able to capture the joint probability distribution accurately ([21]), but it requires a lot of obscure parameters that, in the absence of data, make eliciting expert judgment infeasible ([26]). Another issue with discrete networks is the specification of conditional probability tables. The size of these tables grows with the number of parent nodes and the number of states the parent nodes can take on. Often, the conditional probability tables can become so large that accurately learning the conditional probabilities from data becomes impossible. Furthermore, in reliability analyses, data is often scarce, as failures do not often occur, and eliciting experts becomes fruitless when the conditional probability tables are large ([26]). In 2008 the non-parametric Bayesian network was developed ([27]), which allows variables

to be represented by arbitrary continuous or discrete marginal distributions, and correlations to be specified by (conditional) correlation coefficients, rather than conditional probability tables. The research in this dissertation uses this non-parametric BN, which is described in more detail in Section 2.2.

2.2. NON-PARAMETRIC BAYESIAN NETWORK

The choice to use the non-parameteric BN ([27], [26]) in this dissertation was based on the needs of levee reliability calculations. We often have marginal distributions of the random load and resistance variables that play a role in failure (e.g. soil properties or river water levels), and we often have some basis for estimating correlation coefficients between variables. In particular, these are usually spatial autocorrelations, as soil parameters are generally random processes. As will be explained in this section, the non-parametric BN is very well suited to this type of data availability, whereas traditional, or discrete, BNs would be overburdened by the number of parameters needed to specify the joint distribution. The following frame gives an example of how even a simple network of only 20 spatially autocorrelated nodes can overburden a discrete network.

Consider an example of a soil parameter that is important to the piping failure mechanism, k, which is the permeability of the sand layer under a levee. Suppose we have estimated the marginal probability density using measurements, and a parameter d_x, the correlation length, which describes the strength of the spatial autocorrelation. Let us consider the autocorrelation function in Eq. 2.3, which describes the correlation of a variable with itself at another location in space. It depends on the longitudinal distance between two points (Δx), and the parameter d_x.

$$\rho(\Delta x) = \exp\left(\left(-\frac{\Delta x}{d_x}\right)^2\right) \tag{2.3}$$

Now suppose that we are representing a levee by 20 cross sections, each separated by some distance Δx. The network for the spatial distribution of k is captured in Figure 2.2. According to Eq. 2.3, each node will be correlated with every other node in the network.

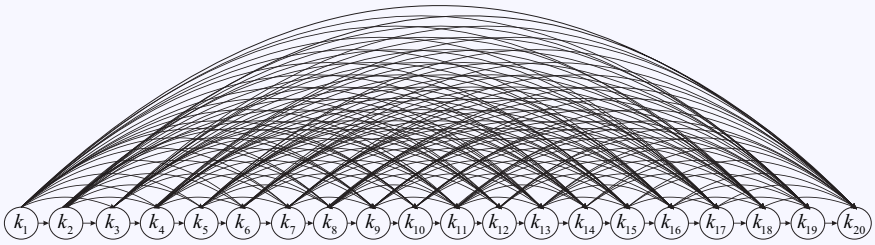

Figure 2.2: BN for the multivariate spatial distribution of k, with 20 cross sections representing a levee segment.

In the network illustrated in Figure 2.2, the root node is k_1, which in a traditional network would be represented by a marginal probability table. The node furthest to the right, k_{20}, has

19 parents ($k_1, ..., k_{19}$). Even if the marginal distribution of k was discretized very coarsely, with only five bins, the conditional probability table for k_{20} would have $5^{20} = 100$ trillion entries. Such a table is infeasible to parameterize, not to mention inaccurate due to the coarse discretization of the marginal distributions.

The non-parametric BN ([28], [29],[30]) was developed to address some of the shortcomings in traditional networks. The name 'non-parametric' is meant to emphasize the fact that no parametric form of the joint distribution is necessary. A good comparison with other hybrid networks, as well as recent applications using the non-parametric hybrid BN, are provided in [26]. Specific attributes that are useful for levee reliability are that the non-parametric BN describes nodes in the network with marginal (continuous or discrete) distributions, specifies arcs in the BN with correlation coefficients, and calculates the dependence structure among the variables using copulas.

Copulas were first introduced by Sklar ([31], ([32])) as a convenient way to build multivariate probability distributions, because they separate the dependence structure from the marginal distributions. The word "Copula" means "link" in Latin, and copulas literally link the marginal distributions together to form the joint distribution. Consider a random vector $X = (X_1, ..., X_n)$, with marginal distribution functions $F_1, ..., F_n$, and a joint distribution function $F_{1,...,n}$. A copula C is a joint distribution function that operates on uniform random variables, and satisfies Eq. 2.4.

$$F_{1,...,n}(X_1, \ldots, X_n) = C(F_1(X_1), \ldots, F_n(X_n)) \tag{2.4}$$

There are many popular copulas, which differ most notably in how they describe tail correlation (see [33] and [34]). The choice of copula is usually determined by observing the tail dependence in data. The non-parametric BN can theoretically take any copula to represent the dependence structure, but using the Gaussian copula makes performing inference more efficient. This is because the Gaussian copula inherits most of the properties of the Gaussian distribution, which in turn allows for analytical derivations of any conditional distributions. In the reliability modeling described in this chapter, the Gaussian copula is used to describe the spatial autocorrelation of the resistance variables. While not proven with data, this description of the correlation is expected to be reasonable, because it does not impose any extra correlation in the extremes (tails) of the distribution, which is appropriate for spatial autocorrelation (i.e. the distance determines the strength of the correlation, not the extremity of the value of the resistance variable).

In reliability analysis, it is common to use the Nataf or Rosenblatt transformation to describe and sample correlated variables. Recent publications have shown that the classic version of Nataf and the Rosenblatt transformations are equivalent to using the Gaussian copula (see [33] and [35]), which is used this dissertation.

One constraint in using the Gaussian copula is that the correlation matrix be positive definite. In the example in Figure 2.2, this would be guaranteed because the autocorrelation function in Eq. 2.3, which is used to calculate the correlations between nodes, is a valid positive definite correlation function. When correlations are based on expert opinion, it is recommended to elicit conditional rank correlations, which for the Gaussian copula can be transformed into Pearson product-moment correlations, using

recursive formulas described in [36]. The reason for this is that experts may choose any conditional rank correlation between -1 and 1, and the transformation algorithms will always guarantee positive-definiteness. In [29], Hanea showed that conditional copulae, together with the one-dimensional marginal distributions and the conditional independence statements implied by the graph uniquely determine the joint distribution.

2.3. METHODOLOGY

This section describes the methodology to construct, sample, and perform inference in a BN for levee reliability. This method is presented for the case that failure of the levee is described by a formula (such as the piping failure mechanism). BNs can be excellent tools in cases where the failure mechanism is not analytically formulated. However, it would require some preprocessing, and falls outside the scope of this dissertation. Specifically, the geotechnical model describing failure would need to be used to extract the dependence between the input random variables and the output variables (e.g. the limit state function). The latter would then be incorporated within the BN as a non-functional random variable, with arcs and correlations representing the dependence extracted via the geotechnical model (see [37] for an example from a different field). Thereafter, the method as presented in this chapter could be applied.

It is useful to clarify some terminology about spatial scales. A levee *system* refers to a large stretch of levees (typically tens of kilometers or more), within which are numerous levee *segments* (typically in the order of 1 kilometer) that are considered statistically homogeneous. This means that while the random variables (e.g. soil permeability) fluctuate within the segment, the parameters of their probability distribution are constant over the segment. The smallest spatial scale considered is a levee *cross section*. This is a slice of the levee over which the values of the random variables are assumed to be constant. Figure 2.3 illustrates a levee segment and a cross section.

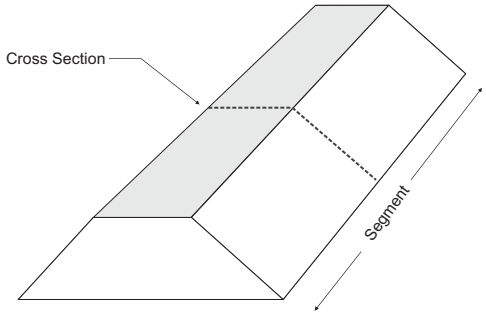

Figure 2.3: Illustration of a levee segment and a cross section.

2.3.1. RELIABILITY OF A LEVEE CROSS SECTION

The method begins by considering the reliability of a cross section. The BN is built on the analytical representation of failure, which is often postulated as a limit state function. Such a function, typically denoted by the letter Z, is positive when the levee is reliable

and negative when the levee fails. A failure node is included in the network, $Fail$, which is 0 when $Z \geq 0$ and 1 when $Z < 0$. As an example, assume that the limit state function depends on two resistance variables: R_1 and R_2, and a load variable S. Figure 2.4 shows what the Bayesian network for the failure probability of the cross section might look like. The resistance and load variables are shown as clear circular nodes, representing input random variables, and Z and $Fail$ are shown as a circular nodes with black edges, representing functional nodes (i.e. their relationship with their parent nodes is specified by an equation rather than a copula). Note that in this example, the random variables are independent of each other (no arcs between them), but this does not have to be the case.

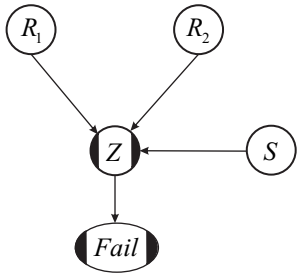

Figure 2.4: Example of a Bayesian network for cross sectional levee failure probability

The probability of failure for the single cross section is described by the integral in Eq. 2.5, which is estimated using the BN. The input variables in the BN are sampled taking into account any defined correlations between variables (see section 2.2 for details). The functional nodes are then calculated using the samples of the input variables. The failure probability estimate $\hat{P}_f$ is estimated according to Equation 2.6, where N is the number of MC samples, and $fail_j$ is the value of the failure node *Fail* (1 or 0) for the j^{th} sample. In the remainder of the dissertation, I drop the notation $\hat{P}_f$ in favor of P_f, with the understanding that these are estimates.

$$P_f = \int_{Z(R1,R2,S)<0} f_{R_1,R_2,S}(R_1,R_2,S)\,dR_1\,dR_2\,dS \tag{2.5}$$

$$\hat{P}_f = \frac{1}{N}\sum_{j=1}^{N} fail_j \tag{2.6}$$

For a cross section in which the parent nodes (R_1, R_2, and S in our example) are uncorrelated, the BN for the cross section is fairly trivial. However, I consider it illustrative because it serves as a building block for the segment BN.

2.3.2. RELIABILITY OF A LEVEE SEGMENT

Levee segments can be long, typically a few kilometers. The failure probability of a cross section is almost always a poor representation of the failure probability of the entire segment. Specifically, unless the variables contributing to failure are fully spatially-correlated throughout the segment, the cross section failure probability will be an underestimate. So instead of representing the failure probability by a single cross section,

it is represented by multiple cross sections, and take care to account for the spatial autocorrelation of the variables between cross sections. Continuing with the example of two resistance variables, R_1 and R_2, and a load variable S, the segment failure probability is represented by the integral in Eq. 2.7. While similar to the cross-sectional integral, the variables R_1 and R_2 in Eq. 2.5 are now the vectors $\mathbf{R}_1$ and $\mathbf{R}_2$, of length n, where n is the number of cross sections in the segment. For example, $\mathbf{R}_1 = \left[R_1^1, R_1^2, \ldots, R_1^n\right]$. The integral is therefore of dimension $2 \cdot n + 1$, assuming that the load is constant over the segment. In general, the dimension will be $n_R \cdot n + 1$, where n_R is the number of resistance variables in a cross section.

$$P_{f,seg} = \int_{\mathbf{Z}<0} f_{\mathbf{R}_1,\mathbf{R}_2,S}\left(\mathbf{R}_1, \mathbf{R}_2, S\right) d\mathbf{R}_1 d\mathbf{R}_s dS \tag{2.7}$$

Figure 2.5 shows what the BN would look like for a levee segment represented by three cross sections (for the case where the cross-sectional BN is described in Figure 2.4). In the figure, superscripts indicate the cross section. So for example, R_1^2 indicates variable R_1 in the second cross section. $Fail^1$, $Fail^2$, and $Fail^3$ represent the failure nodes for the first, second, and third cross sections, respectively. These cross-sectional failure nodes are then connected to a failure node for the entire segment, $Fail^{Seg}$, a binary node (1 for failure and 0 for non-failure), described in Eq. 2.8.

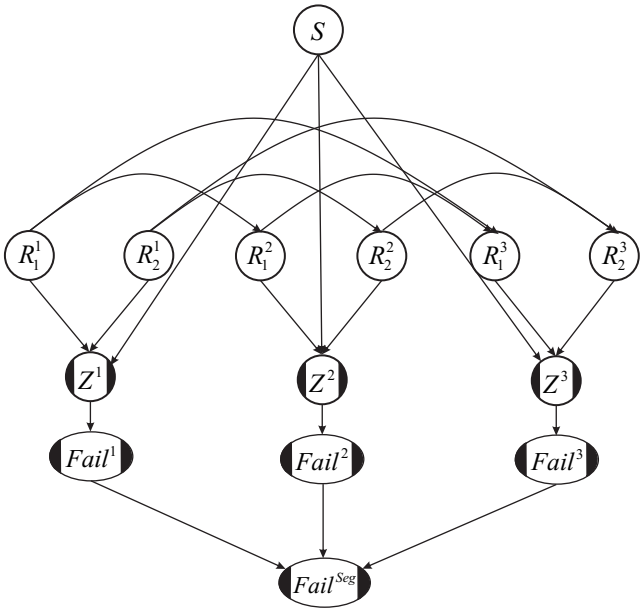

Figure 2.5: BN for a levee segment, in this example represented by three cross sections, each with autocorrelated resistance variables R_1 and R_2, and one common load variable S

$$Fail^{Seg} = \begin{cases} 0, & \text{if } \forall i \ Fail^i = 0 \\ 1, & \text{if } \exists i \text{ s.t. } Fail^i = 1 \end{cases} \tag{2.8}$$

The number of cross sections needed to adequately estimate the failure probability of the segment will depend on the autocorrelation of the resistance variables, the length of the segment, and the magnitude of the prior failure probability. The number of cross sections representing the segment is iteratively increased, each time computing the failure probability of the segment, until additional cross sections no longer change the estimate. The method requires a defined stop criterion, such that when the criterion is met, the number of cross sections is considered sufficient to represent the spatial variability of the segment. In this dissertation, I defined a stop criterion such that the segment failure probability estimates from the previous n_{iter} iterations all lie within the 95% confidence interval of the current estimate. The number of previous iterations to include, n_{iter}, can be based on visual judgment. The specifics are discussed in more detail in Chapter 3.

Arcs between resistance variables are specified with Pearson product moment correlations, which can be estimated using data and one of a number of valid autocorrelation functions ([38]). The one used in this research is commonly used for resistance variables in the Netherlands ([8], [10],[6]), and depends on the distance between variables Δx and the parameter d_x (see Eq. 2.9).

$$\rho(\Delta x) = \exp\left(-\frac{\Delta x^2}{d_x^2}\right) \tag{2.9}$$

Once the marginal distributions of the input random variables have been specified, as well as the equations of the functional variables and the correlation matrix $\Re = \{\rho_{jk}\} = \rho(\Delta x_{jk})$ (see Eq. 2.9), where Δx_{jk} is the distance between R_i^j and R_i^k, the joint distribution over the random input variables can be sampled (described in Section 2.3.4). Note that the correlation function is continuous, but the correlation matrix is discrete because the spatially continuous resistance variables have been chopped up into spatially discrete cross sections. The samples are propagated through the equations for the functional variables in the network, to derive the sample of $Fail^{Seg}$. The failure probability of the system is calculated using the standard MC estimator, which in this case is the mean over the samples of $Fail^{Seg}$.

2.3.3. RELIABILITY OF A LEVEE SYSTEM

Once the number of cross sections that will represent each of the levee segments in our system is determined, the BN of the entire levee system can be built. This essentially consists of connecting the BNs of the segments. Figs. 2.6 and 2.7 show two possible configurations for a system of two segments, each of which are represented by three cross sections. These figures again consider the example of two resistance variables (R_1 and R_2) and one load variable S in each cross section. Note that in real applications, the number of cross sections may be quite large, but it is kept limited here for easy visualization. When building the system network, it is important to identify which variables are correlated between levee segments. In general, levee segments are typically delineated by considering the length over which variables are statistically homogeneous. This often comes down to notable physical attributes, for example a change in stratigraphy. In such cases, it is reasonable to consider resistance variables between segments to be independent. On the other hand, load variables, like the water level in a river, are typically highly

correlated between neighboring segments. They may be constant over the system (Fig. 2.6), or they may be spatially variable, but highly correlated (Fig 2.7).

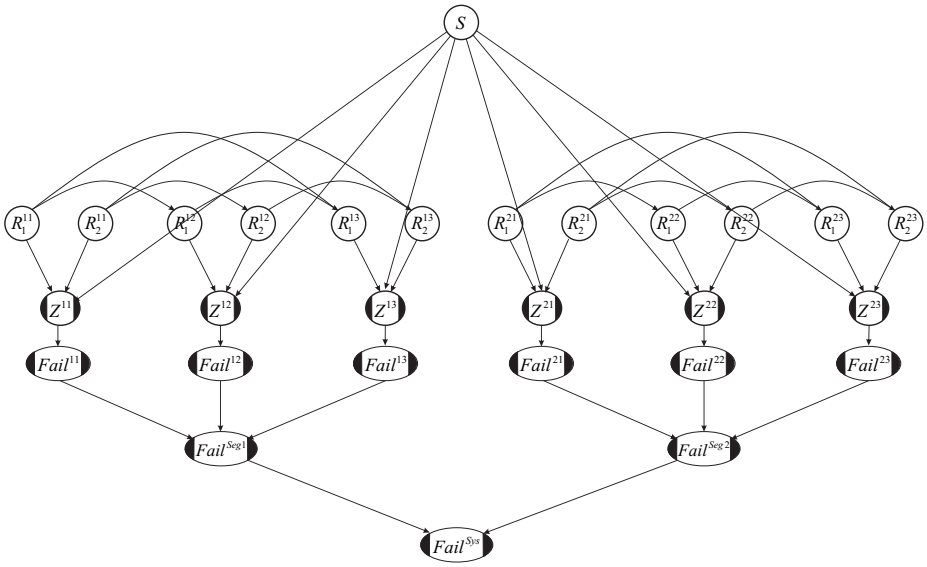

Figure 2.6: Bayesian network for a levee system, in this example represented by two levee segments, each containing three cross sections. The variable S is constant over the length of the system.

2.3.4. SAMPLING

If the resistance variables are autocorrelated, but there is no correlation between the different resistance variables, the spatial distribution of each resistance variable is sampled as follows. For a segment with n cross sections, $[U_1, ...U_n]$ is sampled from the multivariate standard Normal distribution $\Phi_{\Re}(\mathbf{0}, \Re)$, where $\mathbf{0}$ is an $n \times 1$ vector of means equal to zero, and $\Re$ is the $n \times n$ linear correlation matrix, which in the case of the (multivariate) standard Normal distribution is equal to the covariance matrix. To translate these back to the resistance variable of interest, let us consider the example of variable R_1. The variables $[R_1^1, ..., R_1^n]$ are derived using their inverse marginal distributions: $R_1^i = F_{R1}^{-1}(\Phi(U_i))$, $i = 1, ..., n$, where Φ is the standard Normal distribution function.

If there is correlation between variables within a cross section, then this needs to be accounted for in such a way that the entire correlation matrix of the system remains positive definite. There are potentially numerous ways to accomplish this; Appendix G presents a solution derived as part of this dissertation.

2.3.5. INFERENCE USING OBSERVATIONS

Inference is performed differently depending on the type of variable that is observed: an input variable or a functional one. An input variable is described by a marginal probability distribution, whereas a functional variable is described by an equation which operates on the input variables. In the sections below I describe how inference is performed

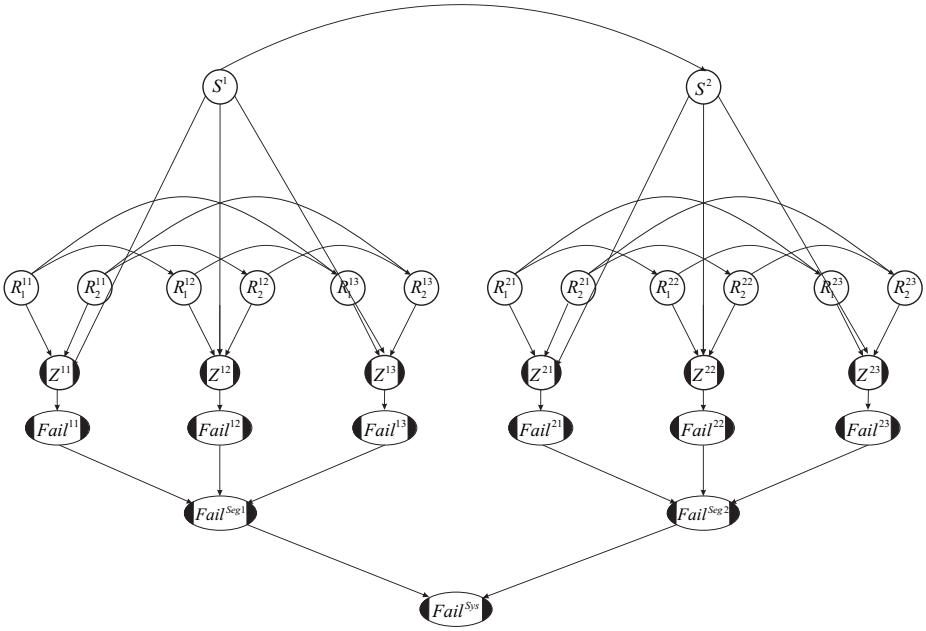

Figure 2.7: Bayesian network for a levee system, in this example represented by two levee segments, each containing three cross sections. The variable S is constant within a segment, and autocorrelated between segments.

for three cases: (1) an observed input variable, (2) an observed functional variable, and (3) a coupled observation of an input and a functional variable (e.g. observed water level and levee survival).

OBSERVED INPUT VARIABLE

When one or more input variables are observed, the conditional joint Gaussian copula (conditional on the observed variable(s)) can be analytically computed. This is straightforward and formulas are available ([27], section 2.4). For ease of reference, the analytical conditioning formulas are provided in the following paragraphs. The ability to analytically condition the joint distribution is the power and benefit of using the Gaussian copula (note that it is also a feature of the multivariate Normal distribution in general). Once the conditional joint copula has been calculated, the marginal distributions of each of the unobserved variables can be used to translate the standard variables back into their original space.

The following discussion provides the formulas necessary to carry our analytical conditioning for the Gaussian copula. Suppose we have a vector of random variables X. This vector is first partitioned into two vectors X_a and X_b: $X = [X_a\ X_b]^{\mathrm{T}}$, where X_a contains the unobserved variables, and X_b contains the observed (conditioning) variables. The means are then partitioned, $\mu = \left[\mu_a\ \mu_b\right]^T$ as well as the covariance matrix, which in the case of standard normal variables is equal to the correlation matrix:

$$\Re = \begin{bmatrix} \Re_{aa} & \Re_{ab} \\ \Re_{ba} & \Re_{bb} \end{bmatrix} \tag{2.10}$$

The conditional joint distribution of X_a given the observed values of X_b (denoted X_b^{Obs}) is then normally distributed with mean $\bar{\mu}$ and covariance matrix $\bar{\Re}$, which are computed according to ([39]):

$$\bar{\mu} = \mu_{\mathbf{a}} + \Re_{ab}\Re_{bb}^{-1}\left(X_b^{Obs} - \mu_b\right) \tag{2.11}$$

$$\bar{\Re} = \Re_{aa} - \Re_{ab}\Re_{bb}^{-1}\Re_{ba} \tag{2.12}$$

In the BN methodology, the (conditional) joint is constructed using standard normal variables U instead of the real-world variables X. The real observations (X_b^{Obs}) are translated to the associated standard normal 'observations' (U_b^{Obs}) via:

$$u_{b,i}^{Obs} = \Phi^{-1}\left(F_{X_{b,i}}\left(x_{b,i}^{Obs}\right)\right) \tag{2.13}$$

Once the conditional joint is derived and sampled, the samples of the unobserved variables are transformed via the marginal distributions of each variable according to Eq. 2.14.

$$X_{a,i} = F_{X_{a,i}}^{-1}\left(\Phi\left(U_{a,i}\right)\right) \tag{2.14}$$

In the case that there is a constant observed load over the segment, and resistance variables that are not dependent on the load (this is generally always the case), the process simplifies because the load can simply be treated as a constant, and the resistance variables can be sampled as they would have without a load observation.

OBSERVED FUNCTIONAL VARIABLE

When a functional variable is observed, the network is first sampled, including the functional variables. This generates an empirical joint distribution over the random and functional variables. The observation can then be imposed as a constraint on the samples. For example, suppose the limit state function is observed to be greater than zero (indicating no failure). The joint samples of all our random variables for which $Z > 0$ would then be retained; these samples would serve as an empirical conditional joint distribution. This is also known as rejection sampling, because all samples for which our condition ($Z > 0$) is not met are rejected. When the variance of the posterior failure probability estimate is too high using rejection sampling, other methods are available, such as importance resampling ([40]), or Markov-chain Monte Carlo ([41]), but those are not considered in this dissertation.

COUPLED OBSERVATION OF INPUT AND FUNCTIONAL VARIABLES

Coupled observations of input and functional variables are often of interest. Most notably in levee system reliability, there is interest in survival observations: water level observations and survival of the levee. These coupled observations make it possible to

update the failure probability estimate, as well as obtain useful information about the remaining uncertain variables in the network.

In the case of coupled input and functional variables, the first step is to (analytically) specify the conditional joint distribution, given the observed value of the input variable. The next step is to sample the conditional joint network, and retain only those samples that meet the observed value(s) of the functional variable (e.g. $Z > 0$ for survival). These retained samples form the updated empirical posterior joint distribution over the resistance variables.

Reducible vs. irreducible uncertainty

The primary purpose of a BN is to reduce uncertainty in a probabilistic model by using observations. However, not all uncertainty can be reduced. Uncertainty in resistance (soil) variables is largely due to shortage of knowledge, and is referred to as epistemic uncertainty. This type of uncertainty is considered fully reducible; that means that given infinite resources, it would be possible to know the value of the soil variables exactly. By contrast, load variables vary in time. Even with infinite measurement capabilities, one cannot know exactly what the water level in a river will be next month. That type of uncertainty, which is referred to as aleatory uncertainty, is irreducible. The BN that is developed and applied in this dissertation focuses on reducing the uncertainty in the resistance variables, with the assumption that the uncertainty in their distribution is entirely reducible. This is a reasonable assumption for soil variables, because the epistemic component of their uncertainty generally dominates the total uncertainty. Thus, the impact of neglecting aleatory uncertainty in resistance variables is probably small, but it should be noted that it is not a conservative simplification. That means that neglecting aleatory uncertainty in resistance variables results in a lower posterior failure probability than when it is accounted for. For the load variables, the uncertainty is generally considered aleatory and thus irreducible. However, the temporal variability is specified with a probability distribution, the parameters of which are typically fitted using limited data. In this sense, load variables also contain epistemic uncertainty (parameter uncertainty), which can be reduced as additional load measurements become available. The BN used in this dissertation does not explicitly model this uncertainty, but could be expanded to do so.

Posterior dependence

Observing the value of a variable in the network can introduce dependence between previously independent variables, which follows directly from the graphical structure according to the concept of D-separation, described in Appendix D. As can be seen from Figure 2.4, R_1 and R_2 are marginally D-separated. That is, R_1 and R_2 are independent, $R_1 \perp R_2$, when no other nodes in the network are observed. Once Z is observed, however, they are D-connected, which means they are conditionally dependent, $R_1 \not\perp R_2 \mid Z$. This dependence is not always easily captured with a correlation coefficient. For example, consider a case where the limit state function is $Z = R_1 + R_2 - S$. Suppose the load S is observed, $s = s_{obs}$, and the levee survived (i.e. $Z > 0$). This means that $R_1 + R_2 > s_{obs}$. Once a value of R_1 is sampled, R_2 is constrained such that $R_2 < s_{obs} - R_1$. Figure 2.8 illustrates the constrained posterior dependence between R_1 and R_2 for the example that $S_{obs} = 12$. For future analysis with the posterior distribution of the resistance variables,

care must be taken not to violate the constrained posterior dependence between them. The simplest way to impose this relationship is to retain the joint resistance variables samples for which $Z > 0$. For example, in Figure 2.8, the posterior samples (shown as black dots) would be retained for future analysis.

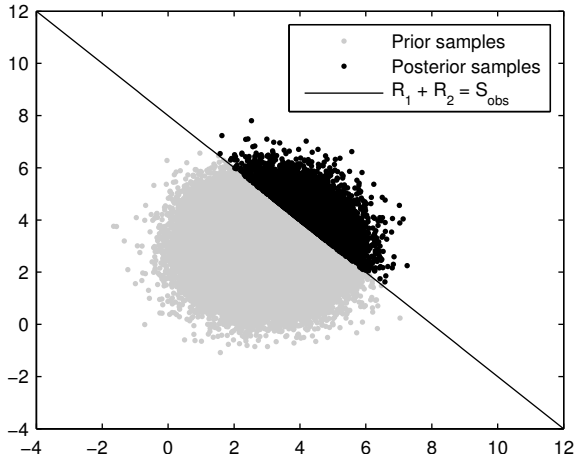

Figure 2.8: Posterior constraints on R_1 and R_2, imposed by the observation that $Z > 0$ for an observed value of S, $S_{obs} = 12$

UPDATED FAILURE PROBABILITY

The posterior segment failure probability $P_{f,seg}^{post}$ can be described by Eq. 2.15. It is the probability that the spatial (multivariate) distribution $\mathbf{Z}$ is less than zero, given that at the time of the load observation (t_{obs}), $\mathbf{Z}$ was greater than zero (survival observed at time t_{obs}).

$$P_{f,seg}^{post} = P\left(\mathbf{Z} < 0 | \mathbf{Z}(t_{obs}) > 0\right) \tag{2.15}$$

Our posterior distribution of the resistance variables includes the condition $\mathbf{Z}(t_{obs}) > 0$, because only joint samples for which this is the case were retained. It is therefore only necessary to calculate the probability that $Z < 0$ using our joint posterior resistance samples. The load is sampled N_p times (where N_p is the number of posterior resistance samples), and calculate Z for each sample. The load varies temporally; therefore observing it at time t_{obs} gives no information about the load at any other time (outside of the temporally correlated window around t_{obs}). This is the reason it must be resampled to calculate the posterior failure probability. The posterior failure probability is then calculated according to Eq. 2.16[2].

$$P_{f,seg}^{post} = \frac{1}{N_P} \sum_{i=1}^{N_P} I\left(Z < 0\right) \tag{2.16}$$

[2]The Z in Eq. 2.16 is computed with the joint posterior resistance samples

The posterior failure probability is calculated using the samples retained after performing inference. If the number of samples is insufficient to keep the variance in the failure probability estimate low (this generally happens when the posterior failure probability is small), Importance Sampling is used, which is a method to reduce the variance in a MC estimate, to sample the load.

Importance Sampling replaces the real load distribution $f_S(S)$ with a biased one $g_S(S)$ that leads to a higher number of failures. The Monte Carlo output is weighted to correct for the use of the biased distribution $g_S(S)$ so that the failure probability estimate remains unbiased; see Eq. 2.17 for the importance sampling estimator, wherein N is the number of samples, and $\mathrm{I}(\cdot)$ is the indicator function.

$$P_{f,seg} = \frac{1}{N}\sum_{i=1}^{N} \mathrm{I}\left(Fail^{Seg} = 1\right) \frac{f_S(s_i)}{g_S(s_i)} \tag{2.17}$$

The choice of the biased distribution $g_S(S)$ will depend on the problem at hand. In general, for updating with survival observations, a reasonable choice is to translate the distribution $f_S(S)$ so that the mean is centered on the observed (high) load S_{obs}.

2.4. DIFFERENCE BETWEEN THE BN AND MONTE CARLO

The BN methodology described in the previous section is essentially a large MC exercise, in which the Gaussian copula is used to correlate many of the samples. This begs the question: what exactly constitutes a BN? Any BN consists of a number of computation steps, written in programming code, that could be carried out without the graphical representation of the network. In my opinion, the network is essential in that it allows the analyst to clarify the connections between variables in the joint distribution. While not strictly necessary, it would be a handicap to proceed without that representation. The BN not only clarifies the connections (correlations) within the joint distribution to the probabilistic analyst, but also to anybody that analyst must communicate with. Errors in probabilistic set up can be caught by experts without a probabilistic background. For example, a geotechnical expert may disagree about a variable being considered constant over a segment; they may argue that while it varies slowly, it will vary somewhat over the length of a segment. The expert has the opportunity to input his knowledge, because the network representation makes it obvious if a variable is constant over a segment. I would argue that without such a network, the geotechnical expert (without a probabilistic background) would never have understood the problem clearly enough to realize this mistake was being made. This is one example, but other errors are also possible: uncorrelated variables that should be correlated, or variables that are not present in the representation that should be there (e.g. model uncertainty). Another advantage of the network representation is that the dependence/independence among variables can be read from the structure of the graph, both before and after one or more of the variables in the network have been observed. Posterior dependence relationships will differ from the prior; that is, variables which were independent prior to observations may become correlated once those observations are included, and vice versa. The concept of D-separation, which is explained in detail in Appendix D, allows one to read these prior and posterior independence/dependence relationships from the graphical

structure. Understanding such relationships without a network, particularly after the inclusion of observations, will become cumbersome and prone to error, especially for large networks. For these three reasons, that the network (1) clarifies the problem to the analyst, (2) allows non-experts in probability theory to contribute to or verify the structure of the network, and (3) allows the analyst to clearly determine prior and posterior independence/dependence among variables in the network, I find the network representation essential, and highly valuable. In the remainder of the dissertation, I refer to the methodology as the BN methodology.

2.5. CONCLUSIONS

In this chapter, I have provided background about different types of BNs, specifics about the non-parametric BN used in this dissertation, and a detailed methodology for describing a spatially extensive levee system using the BN. I described how the network is built for a cross section, for a statistically homogeneous segment, and for a system. I described how the network is sampled, accounting for correlations between variables, to derive samples of the multivariate (spatial) distribution over all the variables in the network, and how the cross-sectional, segment, and system failure probability are calculated. I described how inference is performed in the network based on different types of observations, the potential need for importance sampling in the estimation of the posterior failure probability, and the posterior dependence of the resistance variables. Chapters 3 and 5 show numerical examples and a real-world application, respectively, of the methodology presented in this chapter.

3

THE LENGTH EFFECT AND BAYESIAN UPDATING: VERIFICATION OF THE MODIFIED OUTCROSSING METHOD[1]

3.1. INTRODUCTION

The length effect was first brought to light by Leonardo de Vinci, who said “Among cords of equal thickness the longest is the least strong" ([42]). In the context of levees, where the resistance is spatially variable, the length effect refers to the fact that as the length increases, there is a larger distance over which to encounter a weak spot in the levee, and thus a higher probability of failure.

Different approaches of accounting for the length effect in levees can be found in the literature. Vanmarcke proposed a method involving first crossings ([43], [44]). A crossing refers to the resistance being surpassing by the load, or equivalently the difference between them (also referred to as the limit state function) crossing zero and becoming negative. The method calculates the probability of such a crossing along a given length. More pragmatic methods are also found in the literature. Bowles et al. ([45]) took the length effect into account in a risk analysis of the Herbert Hoover Dike in Florida. In that case, they broke up the levee into segments of about 500 m, and treated them as components in a series system. For all failure mechanisms besides ‘piping through the foundation’, they assumed the segments were independent. For piping through the foundation, they judged that there was ‘some correlation’ between sections, and accounted for this by taking the average of the failure probability assuming (1) full correlation and (2) com-

[1] Parts of this chapter have been submitted for publication as: Roscoe, K., Hanea, A., and Vrouwenvelder, T. (2016). Levee system reliability modeling: The length effect and Bayesian updating. *Submitted to Structural Safety, 2016.*

plete independence. This approach does not clarify for which correlation such an average is valid, or how likely that correlation is to be the correct one. The risk methodology manual developed by the U.S. Bureau of reclamation ([46]) contains guidance for the length effect, proposed by the U.S. Army Corps of Engineers. They break up segments into 'characteristic lengths' which can be considered statistically independent. They do not specify how to estimate the characteristic length other than stating that it can be based on statistical analysis of spatial correlations, or via expert judgment.

In the Netherlands, a modified version of the outcrossing method is used to calculate the length effect, which is referred to in this dissertation as the modified outcrossing (MO) method. It is programmed into the flood defense reliability model ([8], [10]) that is used in national flood risk studies ([47], [6]) and to support the assessment of flood defenses. Reliability calculations are often concerned with limit state functions, which are defined to be negative when geotechnical failure occurs, and positive otherwise. The MO method approximates the limit state function as a one-dimensional Gaussian random field. It is similar to the outcrossing method of Vanmarcke ([43], [44]), but is modified to handle the non-ergodicity of the limit state function. Limited information about the MO method can be found in [8], [10], and [48]; a detailed description is provided in Appendix E. Parallel research is looking into theoretical details of the MO method and its effect on design codes for flood defense systems ([49]). One of the issues addressed in this paper is that although the MO method is an integral part of national flood risk analysis in the Netherlands, its accuracy has not been tested. It approximates the limit state function as a Gaussian random field, but this is generally not the case because it is an (often non-linear) combination of resistance and load variables that are traditionally not Normally distributed).

In this chapter, the BN methodology proposed in Chapter 2 is applied to compute the length effect in synthetic examples. The BN is used to update the joint distribution, and the reliability estimate, using survival observations. These are coupled observations of a (high) water level and survival of the levee, which are available in abundance. Recent research has looked at updating reliability estimates at a cross-section scale ([50]). This is expanded upon in the current research by using the BN to update the reliability of a (long) levee segment.

This chapter has two main objectives: the first is to use the proposed BN methodology to compute the length effect in levee reliability, and the second is to use the BN to address the accuracy/validity of the MO method, both with and without reliability updating using survival observations. The computational efficiency of the two methods is also compared.

This chapter is laid out as follows. Section 3.2 provides an introduction to the MO method. Section 3.3 presents a numerical example via which I compare the prior and posterior segment failure probability estimates of the BN and MO methods. Section 3.4 discusses the computational costs of both methods, and Section 3.5 presents general conclusions. Two appendices are devoted to the MO method: Appendix E provides a detailed description of the MO method. Appendix F describes how updating is performed with the MO method.

3.2. MODIFIED OUTCROSSING METHOD

This section provides a general introduction to the modified outcrossing (MO) method. A detailed description of the method is provided in Appendix E. The MO method begins with the failure probability of a cross section, $P_{f,CS}$, and estimates the segment survival probability $P_{s,seg}$ according to Eq. 3.1, where L is the length of the levee, β is the reliability index ($\beta = \Phi^{-1}\left(P_{f,CS}\right)$), and $\nu_+(\beta)$ is the rate that the limit state crosses zero (from positive to negative), given the reliability index β.

$$P_{s,seg} = \left(1 - P_{f,CS}\right) \exp\left(-L \cdot \nu_+(\beta)\right) \tag{3.1}$$

In order to analytically calculate $\nu_+(\beta)$, the limit state is approximated as a one-dimensional Gaussian random field, such that $Z = \beta - U$, where U is a standard normally distributed variable. Details are omitted from this section, but are provided in Appendix E. Rather, Figure 3.1 is provided to give an intuitive understanding of the crossing rate.

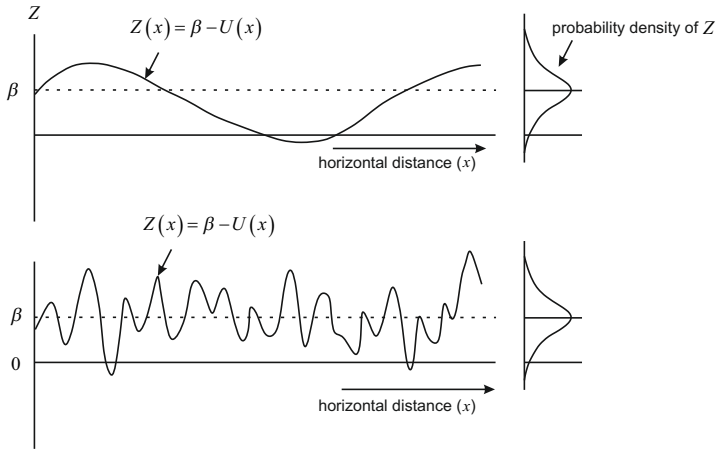

Figure 3.1: Spatial variability of the limit state function Z for a strongly autocorrelated Z function (top) and a weakly autocorrelated Z function (bottom)

Figure 3.1 illustrates the limit state function Z as a random field in one dimension (longitudinally). The probability of having a realization for which $Z < 0$ increases as the length of the segment increases. The increase is dependent on both the length of the levee (L), and how frequently Z crosses 0 (the crossing rate). This latter quantity is dependent on the spatial autocorrelation of Z. For example, as seen in Figure 3.1, a strongly autocorrelated Z function will change slowly in space, while a weakly autocorrelated Z function will show much more rapid change (allowing more opportunities for Z to cross 0).

The crossing rate is calculated analytically based on theory for Gaussian ergodic random fields (see [44]). However, the limit state function is not ergodic, due to the nearly fully-correlated nature of the load over a levee segment[2]. This is taken into account by

[2] Other variables which are fully correlated over the length of the levee segment (such as model uncertainty) also contribute to the non-ergodicity of the limit state function.

calculating the segment failure probability conditional on the non-ergodic part of the limit state function, and then using the theorem of total probability to obtain the full segment failure probability.

3.3. LENGTH EFFECT, NUMERICAL EXAMPLE

This section illustrates and compares the BN and the MO methods via a numerical example. Both prior and posterior failure probability estimates are explored (the latter following from a specified coupled observation of load and levee survival) for levee segment lengths of 500 m, 1000 m, 2000 m, 4000 m, and 6000 m. This section is organized as follows: Section 3.3.1 provides details of the example, Section 3.3.2 describes the criterion used to determine the number of cross sections in the BN, and Sections 3.3.3 and 3.3.4 provide results and discussion about the prior and posterior failure probability estimates.

3.3.1. DETAILS OF THE EXAMPLE

In this example, the cross-sectional limit state function depends on two resistance variables, R_1 and R_2, and a load variable S, with the functional form described by Eq. 3.2. The BNs for a cross section and a segment represented by three cross sections are illustrated in Figures 3.2 and 3.3. Lognormal distributions were assigned to the resistance variables, R_1 and R_2, and a Gumbel distribution to the load variable S. These choices were made to mimic realistic cases, in which resistance variables are commonly described by lognormal distributions, and load variables by extreme value distributions (of which the Gumbel is one). The parameters of the resistance variables are provided in Table 3.1 and the parameters of the load variable in Table 3.2.

$$Z = R_1 + R_2 - S \tag{3.2}$$

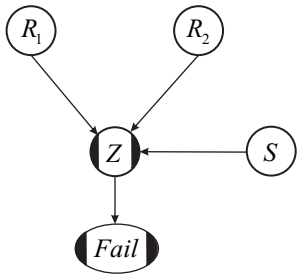

Figure 3.2: Example of a Bayesian network for cross sectional levee failure probability

In this example, I consider survival of the levee for an observed load of $s_{obs} = 4.38$, which corresponds to the 99% quantile of S (i.e. $P(S < s_{obs}) = 0.99$). This seemed a good choice because it is high enough to have an impact on the posterior failure probability, but still low enough to have been realistically observed in measurement records.

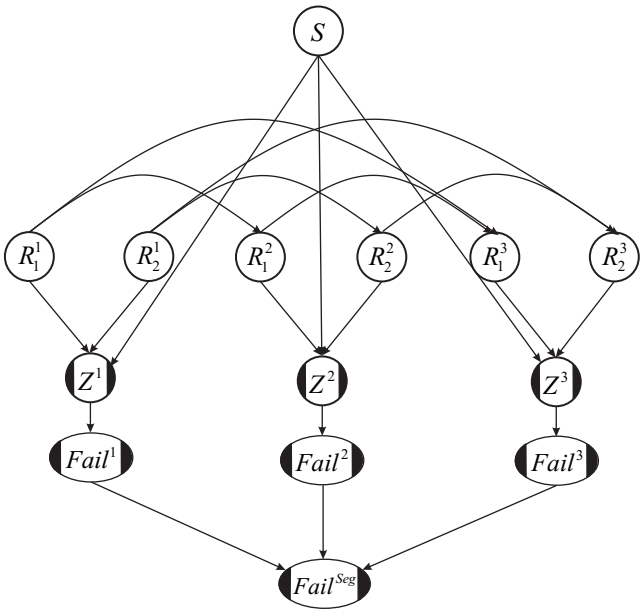

Figure 3.3: BN for a levee segment, in this example represented by three cross sections, each with autocorrelated resistance variables R_1 and R_2, and one common load variable S

Table 3.1: Lognormal distribution parameters μ and σ, distribution mean and standard deviation (SD), and correlation length d_x for resistance variables R_1 and R_2

Variable	Distribution	μ	σ	mean	SD	d_x [m]
R_1	Lognormal	0.842	0.385	2.50	1	200
R_2	Lognormal	1.420	0.232	4.25	1	200

3.3.2. NUMBER OF CROSS SECTIONS IN THE BN

The criterion defined for determining the number of cross sections to sufficiently represent the spatial variability of the segment is based on the width of the 95% confidence interval[3] around the prior segment failure probability estimate, $P_{f,seg}$. The number of cross sections representing the segment were iteratively increased. To speed up the convergence two steps are taken in the iterative process, so that $n = 1, 3, 5$, and so on, computing the segment failure probability estimate each time. The iterative procedure is stopped when eight sequential iterations (e.g. $n = 15$, $n = 17$, ... , $n = 29$) all have esti-

[3] This confidence interval captures the uncertainty due to the variance in Monte Carlo sampling.

Table 3.2: Gumbel parameters μ (location parameter) and β (scale parameter), and correlation length d_x for variable S

Variable	Distribution	μ	β	d_x
S	Gumbel	3	0.3	∞

mates which lie within the 95% confidence interval of the last estimate. At this point, the asymptote is considered to have been reached, so that remaining differences between iterations are due only to sampling variance. The number of iterations for which the estimates must lie within the confidence interval - in our case eight - is somewhat arbitrary, and will require visual inspection of the results to confirm it is a good one.

The confidence interval around the failure probability estimate is computed according to Eq. 3.3, and depends on the relative error ε of the segment failure probability estimate.

$$CI = [(P_{f,seg} - \varepsilon \cdot P_{f,seg}), (P_{f,seg} + \varepsilon \cdot P_{f,seg})] \tag{3.3}$$

The formula for the relative error (see reference [51]) is provided in Eq. 3.4; it depends on the segment failure probability estimate, the number of samples, N, and the value k, which is a quantile of the standard Normal distribution.

$$\varepsilon = \sqrt{\frac{k^2}{N} \frac{\left(1 - P_{f,seg}\right)}{P_{f,seg}}} \tag{3.4}$$

So, for example, since I am interested in 95% confidence intervals, I choose the quantile k such that $\Phi(k) - \Phi(-k) = 0.95$, which is $k = 1.96$.

3.3.3. PRIOR SEGMENT FAILURE PROBABILITIES

The segment failure probability was computed with the BN and the MO method prior to incorporating any survival observations. Figure 3.4 shows (for the 1000 m levee segment) how the BN estimate of the segment failure probability increases with the number of cross sections that represent the segment, and the asymptotic behavior of the estimate once the number of cross sections meets the criterion discussed in Section 3.3.2. The confidence intervals around the BN estimate (shown in Figure 3.4) were calculated according to Eqs. 3.3 and 3.4. The MO estimate is also shown in Figure 3.4 for comparison; it is shown as a horizontal line because it is not a function of the number of cross sections in the BN.

The BN for the 1000 m segment was represented by 41 cross sections. The failure probability estimate is 0.0097 with confidence interval [0.0095, 0.0099]. The MO estimate is 0.0096. The results for the other segment lengths can be summarized by Figure 3.5. The BN and MO estimates are in near-perfect agreement, regardless of the length of the segment. The strong agreement between estimates is a good verification of the MO method.

3.3.4. POSTERIOR SEGMENT FAILURE PROBABILITIES

One of the research questions in this dissertation is how well the MO method approximation holds when a survival observation is taken into account. Recall that details of the implementation of the MO method in combination with updating are provided in Appendix F. The estimates for the failure probability of the 1000 m segment are presented in Figure 3.6. The agreement between the MO and BN methods remains very good. The BN estimates that $P_{f,seg} = 1.59 \cdot 10^{-3}$ and the MO method estimates $P_{f,seg} = 1.63 \cdot 10^{-3}$.

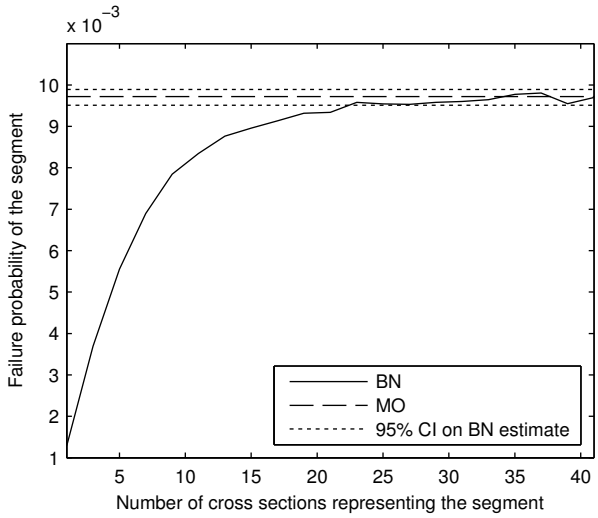

Figure 3.4: Comparison of the BN and MO prior segment failure probability estimates, for $Z = R_1 + R_2 - S$, with R_1 and R_2 lognormally distributed, and S Gumbel-distributed.

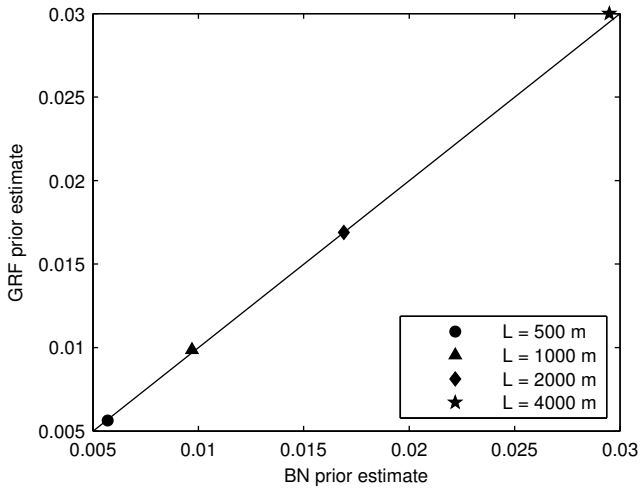

Figure 3.5: Comparison of BN and MO prior segment failure probability estimates (shown together with the 1:1 line, which represents perfect agreement).

Figure 3.7 summarizes the posterior results for levee lengths of 500, 1000, 2000, 4000, and 6000 meters. The differences between the MO and BN posterior segment failure probability estimates remain small, though they increase slightly as the length increases. For a 6000 m segment, the MO method estimates $P_{f,seg} = 5.1 \cdot 10^{-3}$ and the BN estimates $P_{f,seg} = 4.6 \cdot 10^{-3}$, which is a difference of about 10%. This is fairly minor, and in terms of reliability index β (where recall $\beta = \Phi^{-1}\left(1 - P_{f,seg}\right)$), the difference is only 1%.

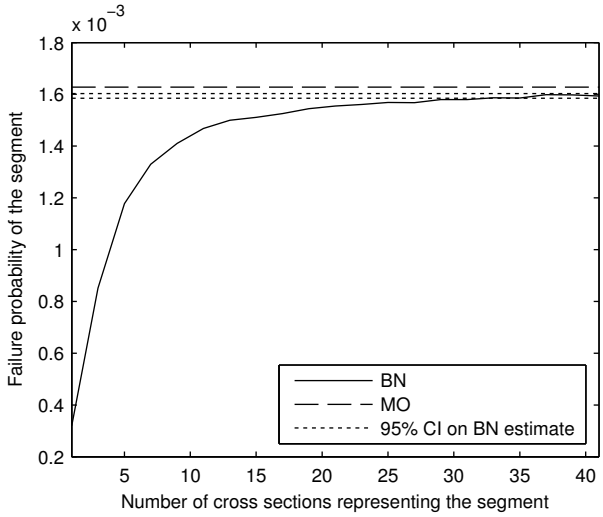

Figure 3.6: Comparison of posterior segment failure probability estimates computed with the BN and MO methods; 95% confidence intervals on the BN estimate are also indicated.

3.4. DISCUSSION

This section discusses the efficiency of the BN and MO methods in terms of computation time[4]. Table 3.3 presents the computation times for the numerical example presented in Section 3.3, for segments lengths of 500, 1000, 2000, 4000, and 6000 m. The calculation time of the MO method does not depend on the length of the segment, and therefore remains relatively constant (fluctuating between 0.5 and 0.7 minutes). The MO method as applied in this chapter included a FORM and a MC analysis at the cross-sectional level. The MC was run to obtain the (more accurate than FORM) reliability index, so that differences between the MO and BN methods were due to the approximations in the MO method, and not to differences in the cross-sectional reliability due to inaccuracies due to the use of FORM. FORM was run to obtain the influence coefficients. Of the 0.5-0.7 minutes that the MO method required, all but a few seconds were spent on the MC calculation of the cross sectional reliability index. The BN requires more time as the number of cross sections needed to represent the segment increases. The MO method is clearly much more efficient, ranging from 6 times faster for shorter segments to 55 times faster for longer segments.

The computation time for the BN is substantially longer than for the MO method, because of the iterative procedure required to determine the number of cross sections (column BN in Table 3.3 includes this iterative procedure). Once the number of cross sections has been determined, the BN is relatively fast (column BN* in Table 3.3), on par with the MO method. Further research can look into more efficient methods to determine the number of cross sections.

[4] Computation times are based on a 2.8 GHz computer with 8GB RAM

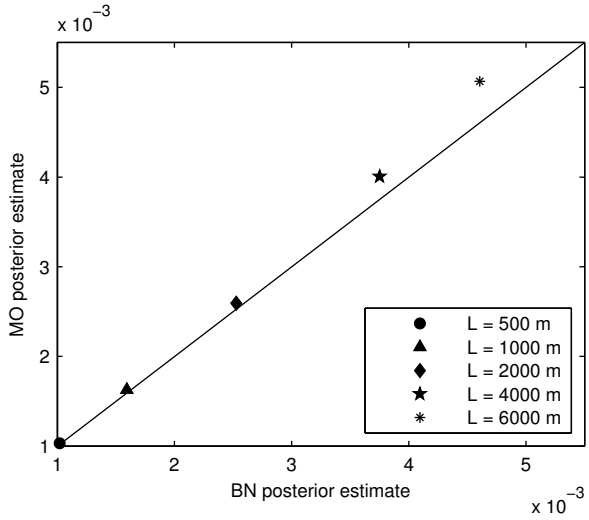

Figure 3.7: Comparison of BN and MO posterior segment failure probability estimates (shown together with the 1:1 line, which represents perfect agreement).

Table 3.3: BN and MO computation times (in minutes) for the example described in Section 3.3, for different segment lengths; BN = BN with iterative procedure to find the number of cross sections (#CS); BN* = BN without iterative procedure.

		Comp. time (mins)		
Length	#CS	BN	BN*	MO
500	43	3.6	0.5	0.6
1000	41	3.9	0.4	0.7
2000	79	16.9	1.3	0.6
4000	81	15	1.0	0.5
6000	111	38.5	2.0	0.7

3.5. CONCLUSIONS

In this chapter, the BN methodology proposed in Chapter 2 was used to calculate the length effect in a levee segment. This is done by sampling the joint spatial distribution of the limit state function, represented by a BN, without having to approximate a parametric form of the spatial distribution. Using Monte Carlo rejection sampling for inference, the method can update failure probabilities of (long) levees using survival observations (i.e. high water levels and no levee failure). Results were compared with the modified outcrossing (MO) method, currently in use in reliability modeling of flood defenses in the Netherlands, via a numerical example, for verification purposes. The primary difference between the two methods is that the BN method samples from the joint spatial distribution, whereas the MO method uses an approximative parametric form of the spatial distribution of the limit state, and solves the problem analytically.

The prior and posterior segment failure probabilities calculated by the two methods

are in strong agreement. Slight discrepancies were found for posterior segment failure probabilities for long segments (4000 and 6000 meters), but these differences were less than 10%, and in terms of reliability index, less than 1%. These results provide a strong verification of the MO method for prior analysis, which is used in the levee reliability model that supports Dutch national flood risk assessments. They also provide an important verification of the MO method for posterior analysis, which has a lot of potential. The speed of the MO method makes it possible to efficiently update failure probabilities of numerous levee segments with abundant survival observations.

Given the strong agreement between BN and MO results, and the relative efficiency of the MO method, I advocate use of the latter in practice. However, the examples considered in this paper do not represent an exhaustive set of cases. For failure probability updating with survival observations, I advocate comparing the BN and MO output for each new type of application (e.g. new limit state function, new set of variable distribution types or correlation parameters). Once the results are verified, the MO method can be used with confidence for all examples of the same type.

Finally, I strongly recommend the use of either the BN or MO method to account for the length effect in reliability analysis over some of the more simplified approaches found in the literature.

4

CORRELATED COMPONENTS IN LEVEE RELIABILITY MODELING: VERIFICATION OF THE EQUIVALENT PLANES METHOD[1]

4.1. INTRODUCTION

System reliability analysis investigates the probability that a system will maintain its functionality; that is, the probability that the system will not fail. Computing the failure probability of complex systems, where the components within the system are correlated, usually requires multi-fold integrals, which are generally impossible to evaluate analytically. Consider a vector of random variables, $\mathbf{X} = [X_1, ..., X_n]$, containing both load and strength variables. The failure of the system is represented by the n-fold integral:

$$P_f = \int_{\Omega(\mathbf{X})} f_{\mathbf{X}}(\mathbf{x})\, d\mathbf{x} \tag{4.1}$$

where $f_{\mathbf{X}}(X)$ is the multivariate density function of X, and $\Omega(\mathbf{X})$ is the failure space, consisting of all realizations of X that lead to failure of the system. The configuration of the failure space depends on how the components in the system are connected: in series, in parallel, or in some hybrid combination. When connected in series, which is typical in levee systems, $\Omega(\mathbf{X}) = \bigcup_i Z_i(\mathbf{X}) < 0$, where $Z_i(\mathbf{X})$ is the limit state function of the i^{th} component, and where failure of each component is defined by $Z_i(\mathbf{X}) < 0$. Monte Carlo methods to estimate the integral in 4.1 are typically prohibitively slow, especially

[1] Parts of this chapter have been published as: Roscoe, K., Diermanse, F., and Vrouwenvelder, T. (2016). System reliability with correlated components: Accuracy of the Equivalent Planes method. *Structural Safety, 57, pp. 53-64.*

in cases where evaluating the limit state functions requires calls to finite element models. A number of methods have emerged in the past decade to address the need for efficient methods to compute system reliability. Sues and Cesare ([52]) proposed a method (Most Probable Point System Simulation, or MPPSS) in which the reliability of the system components is first computed via a method that returns a closed form of the limit state function (e.g. first- or second-order reliability methods). The limit state functions, together with the Boolean expressions defining failure, are then sampled in a Monte Carlo framework. The authors claim that the size of the system is trivial because of the closed form of the limit state functions, but for highly reliable components and/or large systems, it can require billions of samples to acquire the desired accuracy, making this method potentially prohibitively time-consuming. Naess et al. ([53]) proposed a Monte-Carlo-based method in which some tail properties of the distributions are used to substantially improve efficiency. In a follow-up paper ([54]), they tested the method on a large system with thousands of components and found an uncertainty band in which the upper bound is approximately five times the failure probability of the lower bound, for 200,000 samples and a computation time of about 30 minutes to an hour. The method has not yet been tested on systems in which the limit state function requires calls to an intensive external model (e.g. a finite element model), but will most likely be prohibitively slow given the number of samples required. Kang and Song ([55]) proposed an efficient method (sequential compounding method, or SCM) in which the reliability of the components is first computed, and the components are subsequently combined into equivalent components, two at a time, until the full system reliability is obtained. They tested their method on various system configurations, and found very good accuracy for all the configurations considered in the paper. Chun et al. ([56]) presented a complimentary method to SCM, which computes the sensitivity of the system failure probability to the reliability indices of the components. The method does not consider the sensitivity of the system failure probability to the random variables that influence the component reliability indices. In the Netherlands, the reliability of flood defense systems has been a key research area for decades. Based on a series of papers from the 1980s ([57], [58], [59], [60]), an efficient method for combining the failure probabilities of correlated components – referred to here as the Equivalent Planes method – was developed for series systems and implemented in reliability software for the Dutch flood defense system ([8], [10]). The method was designed for series systems (as flood defense systems are primarily connected in series); two components connected in parallel within a system that is primarily connected in series poses no problem, but the method is not intended for systems of numerous components all connected in parallel. Similar to the MPPSS method of Sues and Casare ([52]), the Equivalent Planes method first computes the failure probability of the components, and then replaces their limit state functions with closed-form expressions for subsequent combining. While the MPPSS method allows generic mathematical formulation, the Equivalent Planes method is restricted to linearized forms of the limit state function (hyperplanes). In contrast to the MPPSS method, the Equivalent Planes method does not rely on Monte Carlo methods. Similar to the Sequential Compounding method from Kang and Song ([55]), the Equivalent Planes method combines components sequentially; they differ most notably in the method to derive the correlation between a combined component and the remaining system components.

To accomplish this, the Equivalent Planes method requires information about the autocorrelation of the underlying random variables contributing to failure; the Sequential Compounding method only requires the correlation between components. The Equivalent Planes method was developed to simultaneously meet two requirements for Dutch flood defense reliability modeling: fast computation for large highly-reliable systems, and the ability to compute influence coefficients of both the random variables and the components. These influence coefficients are critical in Dutch flood defense reliability modeling on two fronts: (1) in deltas, where the flood defense system is subjected to loads fluctuating at different time scales, the influence coefficients are needed to scale the failure probability from the time scale of the highest-fluctuating load to the time scale of interest ([10]), and (2) they give flood defense managers a clear overview which variables, levee segments, or failure mechanisms are contributing the most to the failure probability and require the most attention. In the Netherlands, the results of the method – the failure probability of a system of flood defenses – have been used in national flood risk analysis, on which major decisions about the safety standards of the defenses have been based ([61], [47], [62]). However, the accuracy of the Equivalent Planes method for large systems has never been well investigated. Additionally, although the method is in long-standing use, it remains absent from the literature. This chapter serves thus two purposes. The first is to document the method in the literature, and the second is to set up a suite of academic system configurations which can be used to investigate the accuracy of the method.

The chapter is laid out as follows. The Equivalent Planes method is described in Section 4.2; the source of error in the Equivalent Planes method is discussed in Section 4.3; in Section 4.4 various system configurations are presented to explore error propagation and show the performance of the Equivalent Planes method; tolerable error is discussed in Section 4.5; the application of the BN to some of the system configurations is discussed in Section 4.6, and discussion and conclusions are presented in Section 4.7.

4.2. EQUIVALENT PLANES METHOD

The Equivalent Planes method computes the failure probability (P_f) of a system of two correlated components, and - by applying it iteratively - the failure probability of a system of any number of components. The i^{th} component is described by a limit state function, Z_i ; failure occurs whenever $Z_i < 0$. The method starts with two components, connected in parallel (Equation 4.2) or in series (Equation 4.3). Often these components are correlated; that is, failure of one component will influence the failure probability of the second component.

$$P_f = P(Z_1 < 0 \cap Z_2 < 0) = P(Z_1 < 0) \cdot P(Z_2 < 0 | Z_1 < 0) \tag{4.2}$$

$$P_f = P(Z_1 < 0 \cup Z_2 < 0) = P(Z_1 < 0) + P(Z_2 < 0) - P(Z_1 < 0 \cap Z_2 < 0) \tag{4.3}$$

The strategy of the Equivalent Planes method is to replace the conditional probability $P(Z_2 < 0 | Z_1 < 0)$ with an equivalent marginal distribution $P\left(Z_2' < 0\right)$ which incorporates the condition $Z_1 < 0$ by having a non-zero density only in the failure space of component

1. I will describe how the equivalent marginal distribution is computed, but will first highlight the required information for getting started.

4.2.1. GETTING STARTED

To apply the Equivalent Planes method, the failure probability of each of the individual components must be known, as well as the correlation between component failures. The latter is driven by common variables. For example, consider a levee section along a river with two failure modes – overtopping and internal erosion; the water level in the river will influence the failure probability of both components, creating correlation between them. To compute the correlation between components, information about the variables that cause the correlation is needed: (i) their *autocorrelation* – the correlation between a variable in component 1 and the same variable in component 2 – and (ii) *influence coefficients*, which describe how strongly each variable contributes to failure.

The autocorrelation of the variables can be equal to one in some cases (e.g. variables – like water level – which contribute to different failure modes at the same location will be the same for each failure mode). In other cases (consider soil permeability in two neighboring levee sections), the autocorrelations can be obtained from measurements or from expert opinion, or a combination. To obtain the influence coefficients of the variables, the component failure probabilities are computed using first order reliability method (FORM) (for a description of the FORM method, see [63]). FORM approximates the limit state function as a hyperplane at the design point, with the linearized form shown in Equation 4.4; for component i, the coefficients $\alpha_i = [\alpha_{i1}, ..., \alpha_{in}]$ are the influence coefficients corresponding to a vector of standard normal random variables $\mathbf{U} = [U_{i1}, ..., U_{in}]$; the magnitude of each coefficient indicates the relative influence of each variable on component failure. The random variables are standard normally distributed (they are transformed from their actual marginal distributions via FORM), and the influence coefficients are normalized such that $\sum_{k=1}^{n} \alpha_{ik}^2 = 1$.

$$Z_i = \beta_i - \alpha_{i1} U_{i1} - \alpha_{i2} U_{i2} - ... - \alpha_{in} U_{in} \tag{4.4}$$

The component reliability index, β_i, is related to the component failure probability $P_{f,i}$: $\beta_i = \Phi^{-1}\left(1 - P_{f,i}\right)$, where $\Phi^{-1}(\cdot)$ is the inverse standard normal distribution function.

Once the autocorrelations of the variables and the influence coefficients for each component are known, the correlation between components can be calculated according to Equation 4.5.

$$\rho\left(Z_i, Z_j\right) = \sum_{k=1}^{n} \alpha_{ik} \cdot \alpha_{jk} \cdot \rho_{ijk} \tag{4.5}$$

where ρ_{ijk} is the autocorrelation between U_{ik} and U_{jk}. In the remainder of the paper, the symbol ρ, without subscripts, is used to denote the correlation between components; the symbol ρ_{ac} is used to denote the autocorrelation of the variables.

4.2.2. FAILURE PROBABILITY OF A TWO-COMPONENT SYSTEM

The method starts by expressing the limit state function of each component (Z_i) in terms of a single standard normally distributed variable, U_i (Equations 4.6 and 4.7). Note that

this formulation is equivalent to Equation 4.4. The individual random variables and their influence coefficients were needed for computing the correlation between Z_i and Z_j; once this correlation is known, it is more efficient to use the form given in Equations 4.6 and 4.7.

$$Z_1 = \beta_1 - U_1 \tag{4.6}$$

$$Z_2 = \beta_2 - U_2 \tag{4.7}$$

Because the reliability indices in Eqs. 4.6 and 4.7 are constants, the correlation between Z_1 and Z_2 will be the same as the correlation between the variables U_1 and U_2. Variable U_2 can therefore be written as a function of U_1 and an independent standard normally distributed variable U_2^* (Eq. 4.8). This expression can be substituted in the limit state function for the second component (Equation 4.9), ensuring that the correlation between the two components is preserved and that U_2 is still standard normally distributed.

$$U_2 = \rho \cdot U_1 + \sqrt{1-\rho^2} \cdot U_2^* \tag{4.8}$$

$$Z_2 = \beta_2 - \left(\rho \cdot U_1 + \sqrt{1-\rho^2} \cdot U_2^*\right) \tag{4.9}$$

Because of the simplified form of the Z functions, the condition $Z_1 < 0$ is equivalent to $U_1 > \beta_1$ (see equation 4.6). Therefore, to condition on $Z_1 < 0$, the variable U_1 in Equation 4.9 can simply be replaced with a new variable U_1', which captures the tail of the U_1 density function above β_1. The density function of U_1', and how it relates to the density function of U_1 is illustrated in Figure 4.1.

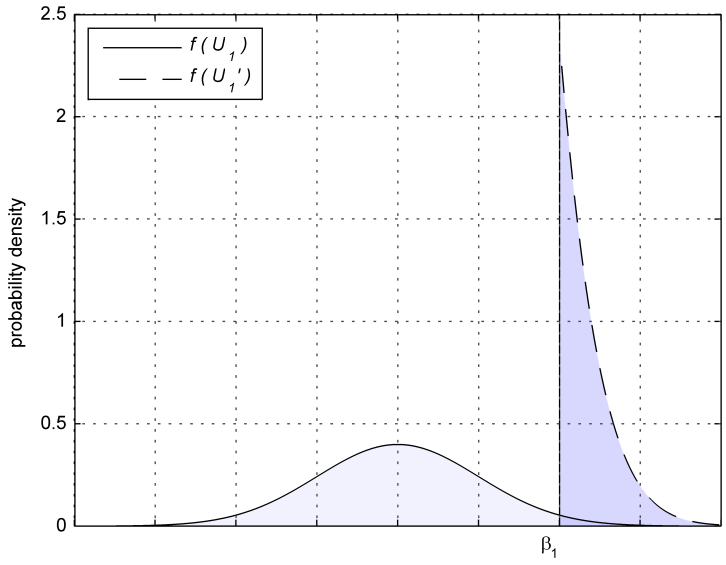

Figure 4.1: Density function of U_1 and U_1'.

The expression for $Z_2^{'}$ is thus:

$$Z_2^{'} = \beta_2 - \left(\rho \cdot U_1^{'} + \sqrt{1-\rho^2} \cdot U_2^{*}\right) \tag{4.10}$$

Note that with Equation 4.10, the problem has been reduced from an n-dimensional problem (where n is the number of variables) to a two-dimensional problem, which is why the Equivalent Planes method is so efficient. Computing the marginal distribution $P\left(Z_2^{'} < 0\right)$ can be done with any probabilistic technique; numerical integration is recommended because $Z_2^{'}$ is only a function of two variables, so even with very small intervals, numerical integration is efficient and very accurate.

Once the conditional distribution $P(Z_2 < 0 | Z_1 < 0)$ has been replaced with the equivalent marginal distribution $P\left(Z_2^{'} < 0\right)$, computing the two-component system failure probability is straightforward; see Equations 4.11 and 4.12 for the parallel and series systems, respectively.

$$P_f = P(Z_1 < 0) \cdot P\left(Z_2^{'} < 0\right) \tag{4.11}$$

$$P_f = P(Z_1 < 0) + P(Z_2 < 0) - P(Z_1 < 0) \cdot P\left(Z_2^{'} < 0\right) \tag{4.12}$$

With the application of Equation 4.11 or 4.12 the failure probability of system of two components can be derived. The next step is to iterate over this procedure to arrive at the multi-component system reliability.

4.2.3. FAILURE PROBABILITY OF A MULTI-COMPONENT SYSTEM

The Equivalent Planes method is iterative, so that once two components have been combined, the two-component system can be considered a new component to combine with a third component, and so on, until all components have been combined. Computing the correlation between the two-component system and a third component presents a challenge. To compute it via Equation 4.5 the influence coefficients for the two-component system are needed; that is, the two-component system needs to be represented by a linearized limit state function (i.e. a hyperplane). Consider Equation 4.4; at the limit state, when $Z_i = 0$, the influence coefficient α_{ik} represents the partial derivative of the reliability index β_i with respect to variable U_{ik}. This can be used to estimate the influence coefficients of the two-component system. In the case where all variables have autocorrelation equal to 1, the influence coefficients can be obtained by numerically estimating $\partial\beta_{ij}/\partial U_k$ (where β_{ij} is the reliability index for the system composed of the two components i and j) for each variable U_k. Note that when the autocorrelation is equal to 1, $U_{ik} = U_{jk} = U_k$. When the autocorrelations are not equal to 1, the concept is similar, but the method to compute the influence coefficients is a bit more complex, because the variable u_k is not exactly the same in component i as it is in component j (i.e. $U_{ik} \neq U_{jk}$). Because U_{ik} and U_{jk} are correlated, one can be written as the function of the other, where the function consists of a correlated ($U_{k,c}$) and uncorrelated part ($U_{k,uc}$). The partial derivatives of the system reliability relative to $U_{k,c}$ and $U_{k,uc}$ are calculated

separately and then combined as shown in Equation 4.13, where α_k^e is the influence coefficient of the k^{th} variable in the equivalent hyperplane for the two-component system.

$$\alpha_k^e = \sqrt{\left(\frac{\partial \beta_{ij}}{\partial U_{k,c}}\right)^2 + \left(\frac{\partial \beta_{ij}}{\partial U_{k,uc}}\right)^2} \tag{4.13}$$

Once the influence coefficients for the combined two-component system have been computed, the correlation between it and a third component can be calculated, and these can be combined these into a three-component system, and so on until all of the components in our system have been combined. Note that method name – Equivalent Planes – comes from expressing a two-component system as an equivalent hyperplane (of the form in Equation 4.4).

4.2.4. PRACTICAL INFORMATION

The methodology described in the preceding sections was programmed in Matlab and has been made freely available via Open Earth Tools (`https://publicwiki.deltares.nl/display/OET/OpenEarth`), which is a repository for free and open source code to handle a variety of problems related to delta and coastal areas ([64]). Open Earth Tools also includes a library of probabilistic tools which are generic and applicable to many problems; the Equivalent Planes algorithm is a part of this library.

4.3. ERROR SOURCE

The Equivalent Planes method is very efficient, but it comes at a price: it is an approximation. This section describes how the approximation introduces error into the system reliability estimate. The focus is on the error incurred using the Equivalent Planes method for combining components with linearized limit state functions. It is important to note that there may also be error introduced in the linearization step; the magnitude of that error is dependent on the behavior of the limit state function, and is not the focus of this chapter. For two components with linearized limit state functions in the form of Equation 4.4, the Equivalent Planes method is exact; error is introduced when a third component is combined with the equivalent two-component system. Figure 4.2[2] illustrates the process by which error is introduced. The method begins with a two-component series system (see Figure 4.2 (a)), with the failure space defined by the area where $Z_1 < 0 \cup Z_2 < 0$, and the original two-component failure probability $Z_1 < 0 \cup Z_2 < 0$. After application of the Equivalent Planes method (see Figure 4.2 (b)), an equivalent limit state function Z_e is obtained, and an equivalent failure space defined by $Z_e < 0$. This step is exact which means:

$$P(Z_e < 0) = P(Z_1 < 0 \cup Z_2 < 0) = P_f \tag{4.14}$$

Figure 4.2 (b) shows the trade that was made in failure space; the area A_1 was released in trade for the area A_2 (see also Table 4.1). This step is exact, so $P(A_1) = P(A_2)$; thus, this can be considered a fair trade. With the introduction of a third component (Figure 4.2

[2] Note that for legibility, Z_i is denoted Zi for $i = 1,2,3,e$ and A_i is denoted Ai for $i = 1,2,3,4$, $Z3$ in Figures 4.2 and 4.3

(c)), the failure space is defined by $Z_e < 0 \cup Z_3 < 0$. Figure 4.2 (c) shows that the fair trade illustrated in Figure 4.2 (b) is now violated. The area A_3, which represents the portion of A_1 that falls in the failure space of Z_3, represents the error in Figure 4.2. This can be explained most clearly as follows.

Consider the failure probability of the original two component system, P_f. If the third component is added to the original two-component system (Figure 4.2 (a)), that will add the area A_{Z3} (shown in Figure 4.2 (c)) to the failure domain. Thus, the system probability becomes:

$$P(Z_1 < 0 \cup Z_2 < 0 \cup Z_3 < 0) = P_f + P(A_{Z3}) \tag{4.15}$$

If the third component is added to the equivalent two-component system (Figure 4.2 (b)), that will add the area A_{Z3} and A_3, and the system probability would be estimated as:

$$P(Z_e < 0 \cup Z_3 < 0) = P_f + P(A_{Z3}) + P(A_3) \tag{4.16}$$

The error that the Equivalent Planes method makes is thus equal to the difference between Eqs. 4.15 and 4.16, which is $P(A_3)$.

Similarly, it can also occur that Z_3 includes some of the gained area (A_2) in its failure space – this area is described as A_4. This situation is illustrated in Figure 4.3. The net error in this case is the probability of A_3 reduced by the probability of A_4.

Table 4.1: Description of the areas in Figures 4.2 and 4.3

Area	Description	Z_1, Z_2	Z_e	Z_3
A_1	Failure space released	$Z_1 < 0 \cup Z_2 < 0$	$Z_e > 0$	$--^+$
A_2	Failure space gained	$Z_1 > 0 \cap Z_2 > 0$	$Z_e < 0$	$--^+$
A_3	Failure space overestimate	$Z_1 < 0 \cup Z_2 < 0$	$Z_e > 0$	$Z_3 < 0$
A_4	Failure space underestimate	$Z_1 > 0 \cap Z_2 > 0$	$Z_e < 0$	$Z_3 < 0$
A_{Z3}	Correct contribution to failure space by	$Z_1 > 0 \cap Z_2 > 0$	$-^*$	$Z_3 < 0$

+ These steps are prior to the inclusion of the third component

* not relevant

4.4. ERROR UNDER VARIOUS SYSTEM CONFIGURATIONS

This section investigates the accuracy of the Equivalent Planes-computed system failure probability estimate for various series system configurations. For this, a *reference calculation* is needed; that is, an estimate of the system failure probability that can be considered exact, with which to compare the Equivalent Planes estimate.

4.4.1. REFERENCE CALCULATION

For systems whose components have equal reliability indices and are equi-correlated, the exact failure probability of the series system can be calculated using the formula:

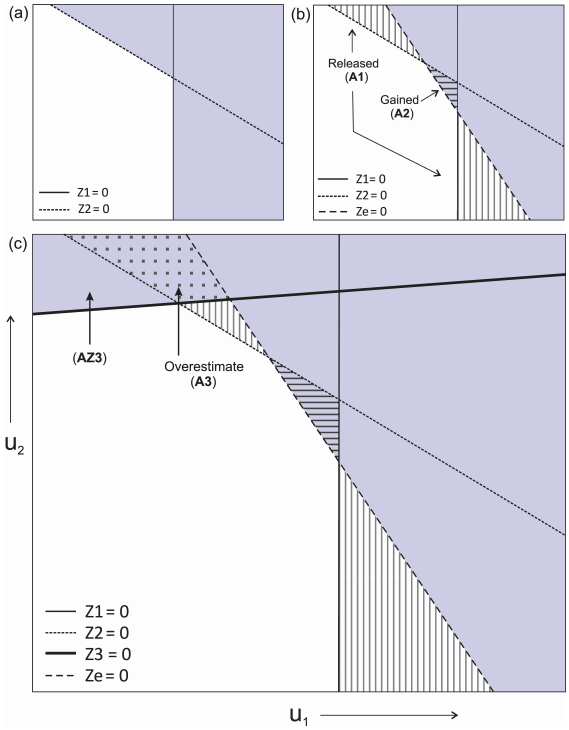

Figure 4.2: Introduction of error in the Equivalent Planes method (overestimate). In each subplot, the blue shaded space represents the failure space. The plots show (a) the original failure space of two components with limit state functions Z1 and Z2; (b) situation after application of the Equivalent Planes method – the equivalent limit state function Ze and the equivalent failure space are shown; A1 is the released area, and A2 is the area that was gained in trade; (c) situation after the inclusion of a third limit state function Z3 – the shaded area is the failure space of Ze and Z3; A3 is the area that was released in (b) in trade for A2, but is recaptured by Z3, violating the fair trade.

$$P = \int_{-\infty}^{\infty} \left\{ 1 - \left[1 - \Phi\left(-\frac{\beta_c - u\sqrt{\rho}}{\sqrt{1-\rho}} \right) \right]^m \right\} \varphi(u)\, du \tag{4.17}$$

where β_c is the reliability index of the components, ρ is the correlation between components, U is a standard normally distributed variable, $\varphi(\cdot)$ is the standard normal density function, and $\Phi(\cdot)$ is the standard normal distribution function.

To compute the 'exact' system failure probability for systems where the components were not equi-correlated, I used a method similar to Sues & Cesare ([52]), only using Monte Carlo Directional Sampling (MCDS) ([65], [66], [67]) instead of crude Monte Carlo. A dynamic sample size criterion was implemented to ensure a high accuracy (described in the following paragraphs). MCDS was considered a good choice because it is relatively efficient compared with crude Monte Carlo, particularly for the case of linearized limit state functions.

In the following sections a brief explanation of directional sampling is provided (in-

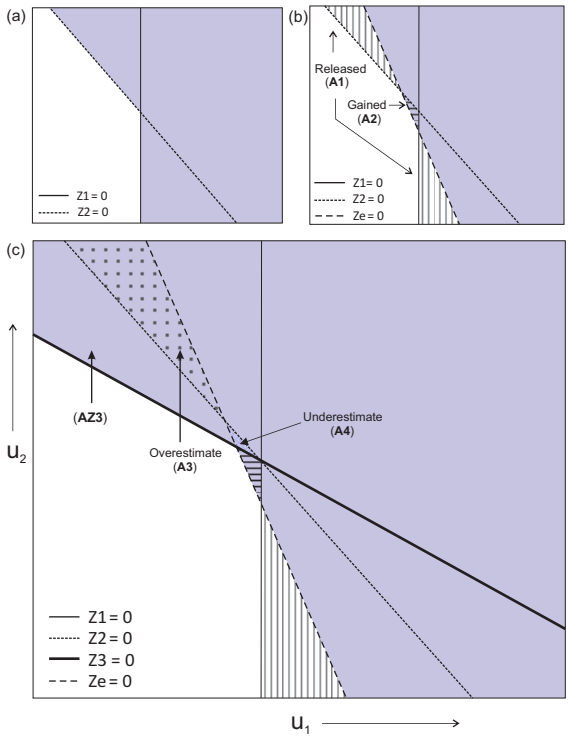

Figure 4.3: Introduction of error in the Equivalent Planes method (over- or underestimate). In each subplot, the blue shaded space represents the failure space. The subplots show (a) the original failure space of two components; (b) situation after application of the Equivalent Planes method – the equivalent limit state function Ze and the equivalent failure space are shown; A1 is the released area, A2 is the gained area; (c) situation after the inclusion of a third limit state function Z3 – the shaded area is the failure space of Ze and Z3; A3 is the area released in (b) in trade for A2, but is recaptured by Z3, the area A4 was gained in (b) but is already part of the system failure due to Z3. A3 and A4 represent a violation of the fair trade in (b).

cluding an efficient approach valid for the case of linearized limit state functions), as well as the implementation of the dynamic sample size criterion.

MONTE CARLO DIRECTIONAL SAMPLING

Directional sampling works by sampling directions in the failure space (Equation 4.18), computing the conditional failure probability given the direction (Equation 4.19), and estimating the failure probability as the mean of the conditional probabilities over all N sampled directions (Equation 4.20). Each direction is defined by a unit vector θ in the standardized normal space; see Equation 4.18. This unit vector is obtained by first sampling all of the n standard normally distributed variables. The vector from the origin to the sampled point in the n-dimensional variable space gives the vector $\mathbf{u}$ in Equation 4.18. The directional unit vector θ is obtained by normalizing the vector $\mathbf{u}$.

$$\theta = \frac{\mathbf{u}}{\|\mathbf{u}\|} = (\bar{u}_1, \bar{u}_2, ..., \bar{u}_n) \tag{4.18}$$

$$P_{f,\theta} = \max\{P_i; P_i = P(Z_i < 0|\theta), i = 1...m;\} \tag{4.19}$$

where m is the number of components in the system.

$$\hat{P}_f = \frac{1}{N}\sum_{j=1}^{N} P_{f,\theta}(j) \tag{4.20}$$

Each direction is defined by a unit vector θ in the standardized normal space; see Equation 4.18. This unit vector is obtained by first sampling all of the n standard normally distributed variables. The vector from the origin to the sampled point in the n-dimensional variable space gives us the vector **u** in Equation 4.18. Normalizing this vector gives us the directional unit vector θ.

In general, Equation 4.19 can be cumbersome to compute because it requires searching in an n-dimensional space for the limit state nearest to the origin (in the given direction). However, the linearity of the limit state functions can be exploited in our case to simplify this process. Let λ be the distance from the origin to the nearest limit state function; then the vector giving the direction and the distance to failure is $\lambda\theta = (\lambda\bar{u}_1, \lambda\bar{u}_2, ..., \lambda\bar{u}_n)$. Because the limit state functions are linear, which means $Z = \beta - \alpha_1 u_1 - ... - \alpha_n u_n$, the failure vector $(\lambda\bar{u}_1, \lambda\bar{u}_2, ..., \lambda\bar{u}_n)$ can be plugged in, and λ can be solved for by setting $Z = 0$.

$$\beta = \alpha_1(\lambda\bar{u}_1) + ... + \alpha_n(\lambda\bar{u}_n) = \lambda\sum_{i=1}^{n}\alpha_i\bar{u}_i \Rightarrow \lambda = \frac{\beta}{\sum_{i=1}^{n}\alpha_i\bar{u}_i} \tag{4.21}$$

I compute λ using Equation 4.21 for each limit state function in our system and take the minimum as our distance to failure. The Chi Squared distribution can then be used to compute the conditional probability given the direction and the squared distance to failure ([67]).

SAMPLE SIZE CRITERION

Our criterion for when the sample size was large enough was a 95% confidence that the difference between the MCDS system reliability estimate ($\hat{\beta}$) and the true value (β) is less than a defined value C:

$$P\left(\left|\hat{\beta} - \beta\right| < C\right) = 95\% \tag{4.22}$$

The important consideration when choosing a value for C in Eq. 4.22 is that it should be small relative to the errors in the Equivalent Planes method, or relative to errors that would be considered important. I chose a value $C = 0.01$, which seemed a good compromise between efficiency (not requiring too many samples) and having an error that was small relative to anything that would be of concern in practice.

The implementation of the stop criterion is described by the flow chart in Figure 4.4. The standard deviation of the failure probability estimate $\hat{P}_f$ is computed as follows:

$$\hat{\sigma}_{P_f} = \sqrt{\frac{1}{N\cdot(N-1)}\cdot\sum_{i=1}^{N}\left(P_{f,\theta}(i) - \hat{P}_f\right)^2} \tag{4.23}$$

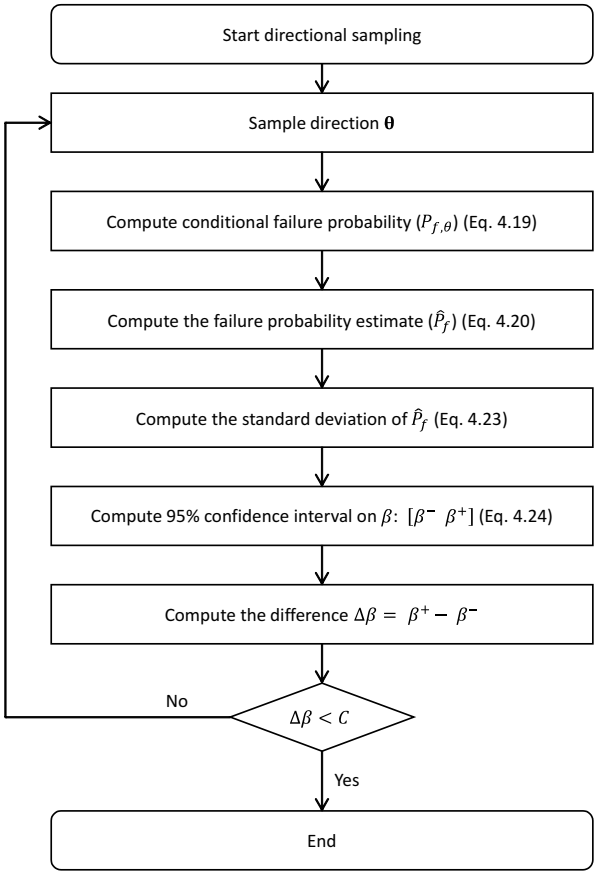

Figure 4.4: Flow chart describing the stop criterion for directional sampling.

The 95% confidence interval on the reliability estimate $\hat{\beta}$ is described by the lower and upper bounds on the interval (β^- and β^+, respectively):

$$\beta^- = -\Phi^{-1}\left(\hat{P}_f + 2\hat{\sigma}_{P_f}\right);\ \ \beta^+ = -\Phi^{-1}\left(\hat{P}_f - 2\hat{\sigma}_{P_f}\right) \tag{4.24}$$

where Φ^{-1} is the inverse standard normal distribution function.

4.4.2. SYSTEM CONFIGURATIONS

I wanted to test various system configurations to explore under which conditions the accuracy of the Equivalent Planes may become a problem. A system configuration was defined based on the following parameters:

(i) n - number of variables

(ii) m - number of components

(iii) ρ - correlation between components

(iv) β_c - component reliability index

(v) ρ_{ac} - auto-correlation of the variables

To test the influence of these parameters a number of system configurations were considered, which are described in the following sections. For each system configuration I computed the 'exact' failure probability (P_f) using either Eq. (4.17) for equi-correlated components, or MC directional sampling otherwise. The exact system failure probability was then compared with the Equivalent Planes-estimated system failure probability ($\hat{P}_f$).

The chosen system configurations were based on a number of considerations. First, I wanted to test extreme system configurations, in an effort to compute bounds on the error in the Equivalent Planes method. To this end, for all the systems considered, the component reliability indices were set to be equal across all the components. This is because if the components of a system have different reliabilities, the smallest component reliability would dominate the system reliability, and would make the errors caused by combining the more reliable components negligible. Second, only series systems are considered, because the Equivalent Planes method was specifically designed with levee systems in mind, which are predominantly series systems. Inclusion of two components connected in parallel within a predominantly series system should not impact the error. However, the method was not designed to compute large parallel systems, and hence such systems were not considered. Third, the most extensively considered cases were with equi-correlated components, because the exact solution can be computed analytically (see Eq. 4.17). This allowed for the investigation of large systems with high reliability indices which would have been too computationally intensive to compute with the MCDS method. Cases where the components were not equally correlated were also investigated, but not as extensively.

CASE I: EQUAL CORRELATION BETWEEN COMPONENTS

Case I investigates series systems with equally reliable components, and equal correlation between all components. The correlation between components was enforced as follows. Assume m limit state functions, each of which is a function of n variables. For a desired correlation between components (ρ), the influence coefficients of the variables were set equal for all m limit state functions ($\alpha_{1i} = \alpha_{2i} = ... = \alpha_{mi};\ i = 1...n$), and set the autocorrelation of all of the random variables equal to ρ. Eq. 4.5 then reduces to:

$$\rho\left(Z_i, Z_j\right) = \sum_{k=1}^{n} \alpha_{ik} \cdot \alpha_{jk} \cdot \rho_{ijk} = \rho \sum_{k=1}^{n} \alpha_k^2 = \rho \tag{4.25}$$

The number of components (m) was also varied, and the reliability index of the components (β_c). The number of variables was fixed to three ($n = 3$). Note that because the variables are only partially autocorrelated, the dimensionality of the problem is much higher than 3; in fact the dimensionality will be equal to the product of n and m (the number of variables and the number of components). Table 4.2 summarizes the system configurations that were considered for the case of equal correlations.

Table 4.2: Parameters defining the system configurations for the Case I

Variable	Value(s)
Number of variables (n)	3
Number of components (m)	3-250
Component reliability (β_c)	3, 4, 5, 6
Autocorrelation (ρ_{ac})	ρ
Correlation between components (ρ)	0.1-0.9 (in increments of 0.1), 0.95, 0.99

The exact system reliability was calculated using Eq. 4.17, but I also ran the MCDS for several of the cases, to assess the computation time involved. In typical cases, Eq. 4.17 cannot be used, and it is then useful to compare the computation time of the MC procedure with the EP method. The computation times are presented in Section 4.7. The results are presented in Figs 4.5-4.8, one figure for each of the correlations 0.2, 0.5, 0.7, and 0.9. It is clear that the Equivalent Planes method performs best for components with high reliability indices, and where the correlation between components is not too high. Figure 4.9 highlights the relationship between the error in the Equivalent Planes method and the correlation between the components; the error is given as a factor difference in the failure probability, which is the ratio of the Equivalent Planes-computed failure probability to the exact failure probability. Table 4.3 shows the factor difference in the failure probability for systems with 250 components. It shows that in the worst case the system failure probability estimated with Equivalent Planes is 2.5 times the exact system failure probability. In the best cases, they are equal. Table 4.4 shows the error in the Equivalent Planes estimate of the reliability indices. The conclusions are the same as for Table 4.3, but viewing the error in terms of reliability index will be useful in the example of how to assess acceptable error in Section 4.5.

Table 4.3: Factor difference in the system failure probabilities (ratio of Equivalent Planes-computed to exact) for the configurations in Case I, for a system with 250 components

	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.7$	$\rho = 0.9$
$\beta_c = 3$	1.3	1.8	2.2	2.2
$\beta_c = 4$	1	1.4	2	2.4
$\beta_c = 5$	1	1.1	1.6	2.5
$\beta_c = 6$	1	1	1.3	2.4

Table 4.4: Error in the Equivalent Planes reliability indexes for the configurations in Case I (difference between Equivalent-planes computed and exact), for a system with 250 components

	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.7$	$\rho = 0.9$
$\beta_c = 3$	-0.18	-0.36	-0.40	-0.33
$\beta_c = 4$	-0.02	-0.13	-0.22	-0.26
$\beta_c = 5$	0.00	-0.03	-0.11	-0.20
$\beta_c = 6$	0.00	-0.01	-0.05	-0.16

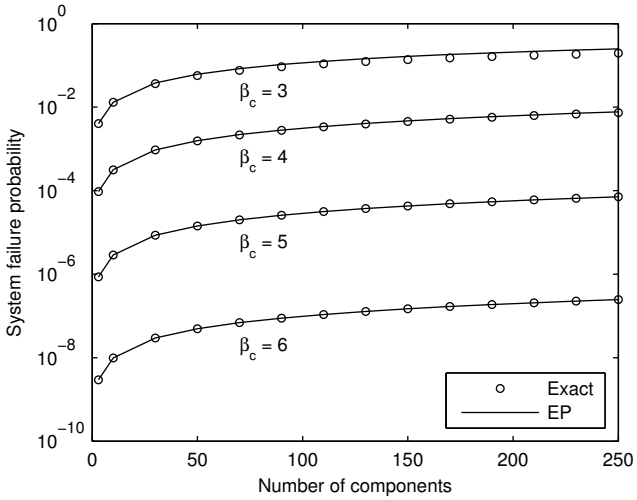

Figure 4.5: Performance of the Equivalent Planes (EP) method for series systems with 3-250 components, all equi-correlated with correlation coefficient 0.2.

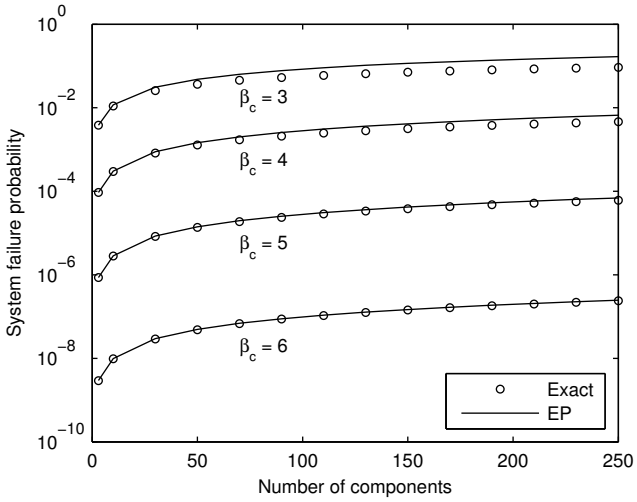

Figure 4.6: Performance of the Equivalent Planes (EP) method for series systems with 3-250 components, all equi-correlated with correlation coefficient 0.5.

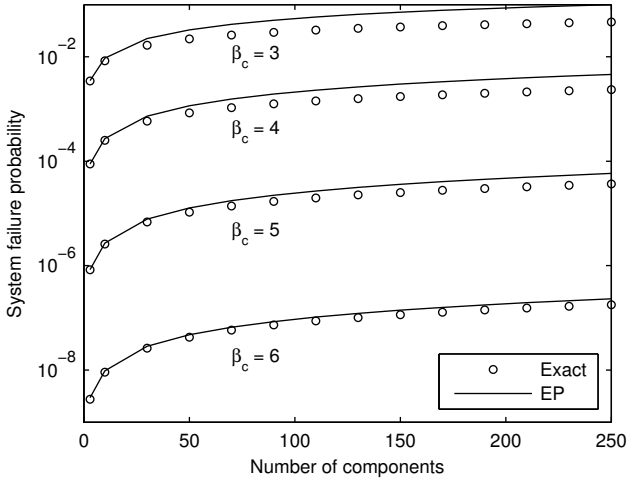

Figure 4.7: Performance of the Equivalent Planes (EP) method for series systems with 3-250 components, all equi-correlated with correlation coefficient 0.7.

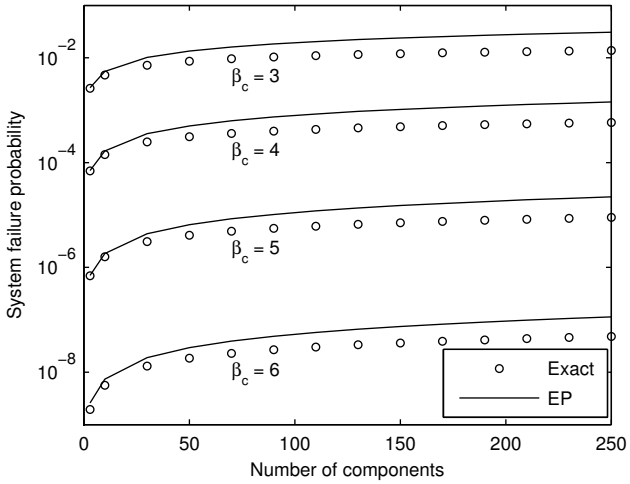

Figure 4.8: Performance of the Equivalent Planes (EP) method for series systems with 3-250 components, all equi-correlated with correlation coefficient 0.9.

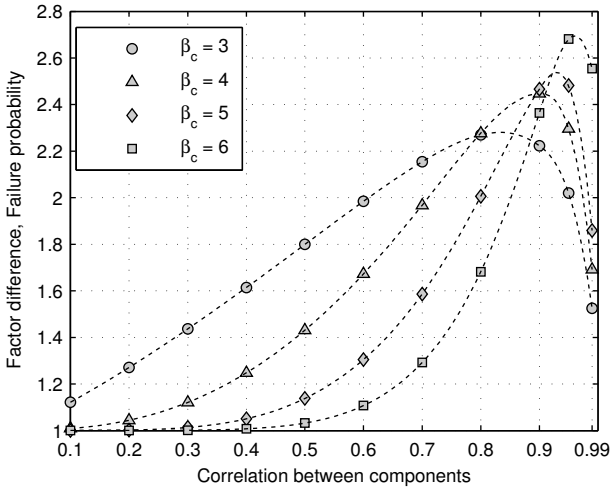

Figure 4.9: Factor difference in system failure probabilities (ratio of Equivalent Planes-computed to exact) for series systems with 250 components, for correlation between components ranging from 0.1 to 0.99.

CASE II: UNEQUAL CORRELATION BETWEEN COMPONENTS

Case II investigates a series system of components with equal reliability indices but unequal correlation coefficients. The correlation between component i and component j in an m-component system was computed according to Eq. 4.26, as in [55]; this formulation ensures the correlation matrix is positive definite.

$$\rho_{ij} = 1 - \frac{|i-j|}{m-1}, \quad i,j = 1,\ldots,m \tag{4.26}$$

Enforcing the correlation structure given by Eq. 4.26 requires choosing the right mixture of influence coefficients and autocorrelations for the random variables in the limit state functions, given the constraint in Eq. 4.5. The number of variables in each limit state function was set equal to m (the number of components), and set the autocorrelation of each variable equal to 1. This reduces the relationship between the influence coefficients and the correlation matrix (see Eq. 4.5) to:

$$\boldsymbol{\rho} = \boldsymbol{\alpha}^T \boldsymbol{\alpha} \tag{4.27}$$

The advantage of having the correlation matrix in the form of Eq. 4.27 is that the influence coefficients for any positive-definite correlation matrix can be derived as the Cholesky decomposition of $\boldsymbol{\rho}$, which is a standard function in packages such as Matlab.

The 'exact' system reliability was computed using Monte Carlo directional sampling. I considered 5-, 10-, and 50-component systems, and component reliability indices of 3, 4, and 5. I did not consider systems larger than 50 components, because of the computational cost of the Monte Carlo simulations. For example, for 50 components and a component reliability index of 5, $1.3 \cdot 10^8$ samples were needed for the MCDS to converge (see sample size criterion in section 4.4.1). The details of this case are given in Table 4.5.

The results are shown in Figure 10. There is very good agreement between the Equivalent Planes-computed system failure probability and the MC-computed system failure probability.

Table 4.5: Parameters defining the system configurations for the Case II

Variable	Value(s)
Number of variables (n)	$n = m$
Number of components (m)	5, 10, 50
Component reliability (β_c)	3, 4, 5
Autocorrelation (ρ_{ac})	1
Correlation between components (ρ)	See Eq. 4.26

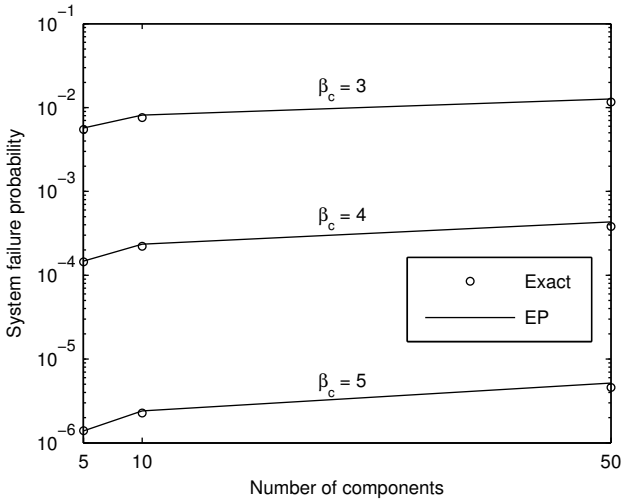

Figure 4.10: Performance of the Equivalent Planes (EP) method for series systems with 5, 10 and 50 components, correlated according to the correlation structure defined in Eq. 4.26.

CASE III: LIMIT STATES WHICH SPAN ALL DIRECTIONS

This case investigates an extreme situation in which the directions of the linearized limit state functions of the components span a three-dimensional space. For many components – all with equal component reliability indices – this begins to enclose a spherical safe region. This is a very unrealistic situation, but is useful for testing how the method performs under such extremes.

To specify the limit state functions, the influence coefficients for each of the three variables needed to be calculated, for every limit state function direction. This was done by generating directional normal vectors (perpendicular to the limit state hyperplane) by using a three-dimensional integer-based grid. A maximum integer was chosen, x_{max}, and a grid was constructed with points placed at all integers from $-x_{max}$ to x_{max}. The directional vectors were generated by connecting these points to the origin. The influence

coefficients can be derived directly from the directional vectors, once they have been normalized (recall the influence coefficients describe a normalized directional vector). Duplicates can arise when the line between the origin and multiple grid points share the same angle (e.g. consider a two-dimensional grid with the points [1, 1] and [2, 2]). The number of limit state functions (i.e. the number of components m) is then a function of the total number of integers (which is equal to $2x_{max}+1$), the number of dimensions n, and the number of duplicates d:

$$m = (2x_{max}+1)^n - d \tag{4.28}$$

An example of a set of directional normal vectors is illustrated for three dimensions in Figure 4.11.

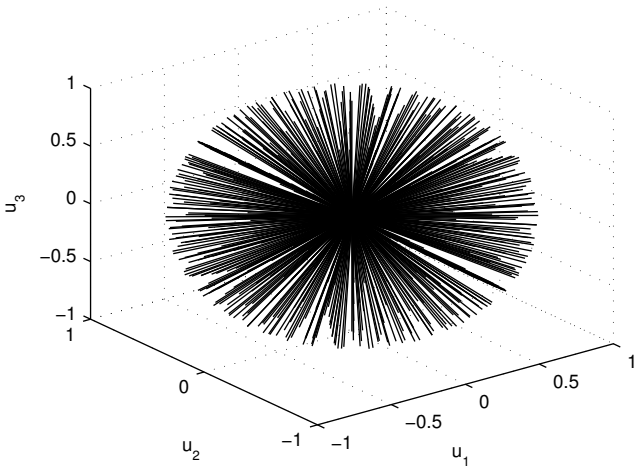

Figure 4.11: Directional normal vectors defining the limit state functions which densely span a 3-dimensional space.

Because the number of components grows exponentially with the number of variables, this case was restricted to three dimensions. An autocorrelation equal to 1 was chosen for all of the variables to limit the dimensionality of the case, in order to make the reference (MC) computations more efficient. The correlations between components (which are variable in this case) were not explicitly chosen, but rather were calculated according to Equation 4.5, based on the influence coefficients. Table 4.6 summarizes the system configurations that were considered.

The results of Case III are presented visually in Figure 4.12. Table 7 shows the factor difference in the system failure probability, and Table 4.8 shows the error in the system reliability index. The worst case tested – 598 components and component reliability indices of 3 – results in an Equivalent Planes system failure probability estimate that is 2.8 times the exact (Monte-Carlo computed) system failure probability. In the best case – 26 components and component reliability indices of 6 – the errors in the Equivalent Planes method are negligible.

Table 4.6: Parameters defining the system configurations for the case spanning an n-dimensional space

Variable	Value(s)
Number of variables (n)	3
Max integer (x_{max})	1, 2, 3, 4
Component reliability (β_c)	3, 4, 5, 6
Autocorrelation (ρ_{ac})	1

Table 4.7: Factor difference in the system failure probabilities (ratio of Equivalent Planes-computed to exact), for the configurations in Case III

	$m = 26$	$m = 98$	$m = 310$	$m = 598$
$\beta_c = 3$	1.3	1.9	2.5	2.8
$\beta_c = 4$	1.2	1.7	2.2	2.6
$\beta_c = 5$	1.1	1.5	2.0	2.4
$\beta_c = 6$	1.0	1.4	1.9	2.2

Table 4.8: Error in the Equivalent Planes system reliability indexes (difference between Equivalent Planes-computed and exact), for the configurations in Case III

	$m = 26$	$m = 98$	$m = 310$	$m = 598$
$\beta_c = 3$	-0.12	-0.30	-0.43	-0.49
$\beta_c = 4$	-0.05	-0.16	-0.25	-0.30
$\beta_c = 5$	-0.02	-0.09	-0.17	-0.21
$\beta_c = 6$	-0.01	-0.06	-0.11	-0.15

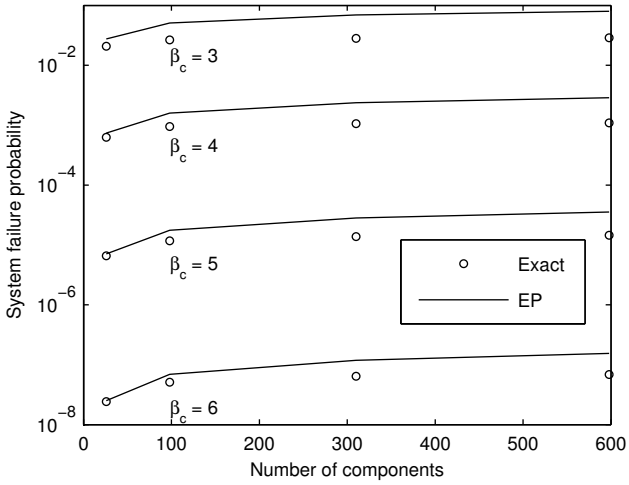

Figure 4.12: Performance of the Equivalent Planes (EP) method for series systems with components whose linearized limit state functions span a 3-dimensional space, with unequal correlation coefficients between components, for 26, 98, 310, and 598 components.

CASE IV: UNCORRELATED COMPONENTS

This case investigates the extreme situation where all components are uncorrelated. This situation is easy to compute analytically (Eq. 4.2 reduces to the product of the component probabilities). It is relevant to test how the method performs under this extreme, since some cases might be nearly uncorrelated in practice. The analytical solution was used for comparison of the results. To set up the uncorrelated case, the number of (independent) random variables was set equal to the number of components ($n = m$), where each component depends on only one of the variables, which will have an influence coefficient of 1 (see Eq. 4.4). Each limit state function is written in terms of all of the variables, but all but one of the influence coefficients will be zero. The number of components was chosen, which then determines the number of variables. The value of the autocorrelation is irrelevant, because each variable appears in only one limit state function. Table 4.9 summarizes the system configurations that were considered for the case of uncorrelated components.

Table 4.9: Parameters defining the system configurations for the case of uncorrelated components

Variable	Value(s)
Number of variables (n)	m
Number of components (m)	3 – 250
Component reliability (β_c)	3, 4, 5, 6
Autocorrelation (ρ_{ac})	Irrelevant (each variable appears in only one component)
Correlation between components (ρ)	0

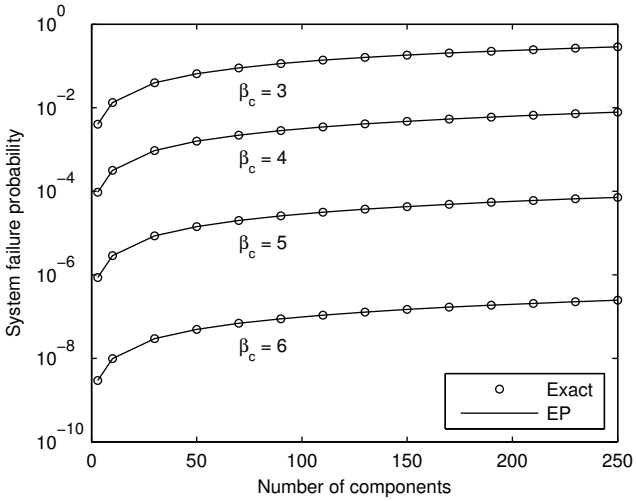

Figure 4.13: Performance of the Equivalent Planes (EP) method for series systems with 3-250 uncorrelated components

4.5. ACCEPTABLE ERROR

This section focuses on translating the error in the Equivalent Planes method to real, tangible terms (e.g. costs), so that acceptable error can be determined. This translation is application-dependent, but I provide an example here for the case of levee systems, principally for illustration purposes. Each specific application will need its own case-specific analysis, but the approach should be similar.

EXAMPLE: LEVEE SYSTEMS

I considered the application to levee systems and provide a simple example of how errors in system reliability can be translated to impacts on levee design, and further to costs. The idea of this section is not to provide full rigor in determining the impact on levee design and costs, but to obtain an order-of-magnitude estimate that can help us decide if the error in our estimate is tolerable.

The analysis was done for the failure mechanism overflow, which essentially considers whether the crest height of the levee is high enough to hold back extreme water levels. The limit state function for overflow is the difference between the levee height (H) and the water level at the levee (W).

$$Z = H - W \tag{4.29}$$

I assigned a Type I generalized extreme value distribution to the water level, with a scale parameter of 0.28 and a location parameter of 2.6, which corresponds to 1/100, 1/1000, and 1/10000 year water levels of 3.9, 4.5, and 5.2 meters, respectively. The system reliability index, β, informs us of the failure probability, or equivalently (in this case) the probability that the water level is higher than the levee height.

$$P_f = \Phi(-\beta) = P(H - W < 0) = P(W > H) \tag{4.30}$$

Because the distribution of the water level F_W is known, the design height – that is, the levee height that corresponds to the reliability index β - can be calculated:

$$H = F_W^{-1}(1 - \Phi(-\beta)) \tag{4.31}$$

Equation 4.31 can be used to translate errors in the system reliability to errors in the design levee height. I considered system reliability indices $\beta = \{3, 4, 5, 6\}$, and errors in the system reliability $\varepsilon = \{0.01, 0.05, 0.1, 0.2, 0.5\}$, and computed the difference in design levee height between the erroneous and true system reliabilities; the results are presented in Table 4.10.

Table 4.10: Difference in design levee height (in meters) due to errors in the system reliability estimate, for different system reliability indexes (β) and different error magnitudes (ε)

	$\varepsilon = 0.01$	$\varepsilon = 0.05$	$\varepsilon = 0.1$	$\varepsilon = 0.2$	$\varepsilon = 0.5$
$\beta_c = 3$	0.01	0.05	0.09	0.19	0.48
$\beta_c = 4$	0.01	0.06	0.12	0.24	0.61
$\beta_c = 5$	0.01	0.07	0.14	0.29	0.75
$\beta_c = 6$	0.02	0.09	0.17	0.34	0.885

To translate the design height differences into costs, I used cost curves derived as part of a national cost-benefit analysis of flood protection measures in the Netherlands ([62],[68], [69]). I considered the extreme cases of a very rural levee and a very urban levee. A very rural levee is one in which the area surrounding the levee is undeveloped, and thus allows for easy expansion of the levee base when the levee is heightened; this ensures that the slope of the levee does not become too steep. Such an expansion will be impossible for a very urban levee because the surrounding area is already fully developed. In such cases, retaining walls are often required to compensate for the steep slope resulting from the levee heightening. Furthermore, in the urban case, a road and a bike lane are also typically present on the levee. For both the rural and urban case there is a base cost – that is, a portion of the cost that is height-independent. In the rural case, this is the cost of removing and replacing the levee revetment; in the urban case, it is the cost of the retaining wall and the road and bike lane. Figure 4.14 shows example cost curves for rural and urban levees. The costs are expressed per km of levee, and as a function of the required levee height increase.

To use the cost curve (Figure 4.14) to translate the error in design height to costs, one must consider the specific case at hand. For example, consider a situation where the current levee must be raised by 1 m to satisfy a required reliability; however, the error in the reliability estimate leads to the belief that it must be raised by 1.5 m. From the curve, the difference in costs between 1 m and 1.5 m can be estimated: approximately 2 million Euros per km of levee for the rural case, or about 4 million Euros per km for the urban case. Whether this error is tolerable depends on a number of factors that are case-specific. Note that the converse situation – that the error would be an underestimate of the levee design height – would result in a less expensive improvement measure, but a

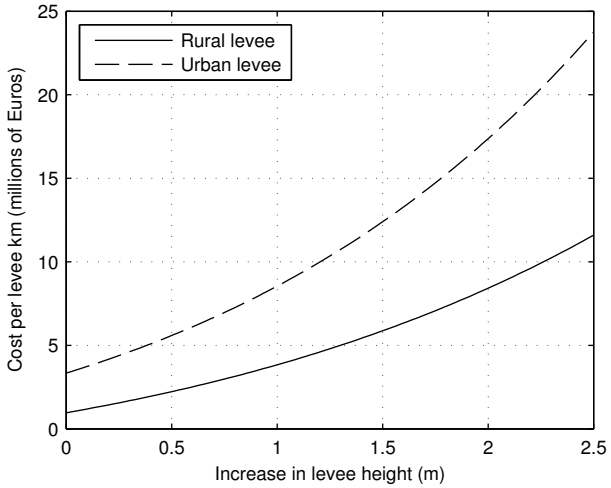

Figure 4.14: Cost curve for levee heightening, for a rural and urban levee, based on data for Dutch levees.

higher risk. The costs associated with the increased risk are more complex to assess than the costs of an improvement measure.

4.6. COMPARISON WITH BAYESIAN NETWORK

In this section I compare how the Equivalent Planes method compares with the BN for each of the cases considered in Section 4.4. For each case we illustrate the BN for the system failure probability, using an example of four components to keep the visualization clear. Each component is represented by its limit state function Z; correlations are represented by arcs between components, and each component feeds into a failure node, $Fail$, which is equal to 1 when $Z_i < 0$ for any component i. The system failure probability is then equal to $P(Fail = 0)$.

For the case of equal correlations, the BN is illustrated in Figure 4.15. The arcs connecting each of the components is specified with a correlation coefficient ρ, which is the same between each pair of components.

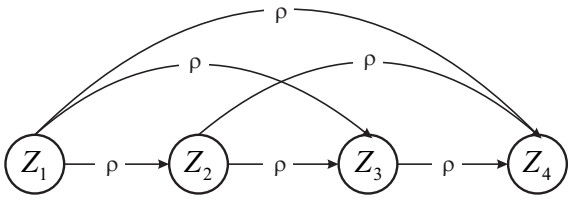

Figure 4.15: BN for the case of equal correlation between components, for an example of four components.

For the two cases of unequal correlations (Case II: Unequal correlation between components, and Case III: Limit states which span all directions), the BN is illustrated by Fig-

ure 4.16. The arcs between components are specified by correlation coefficients, which are calculated according to Equation 4.26 for Case II, and by Equation 4.5 for Case III.

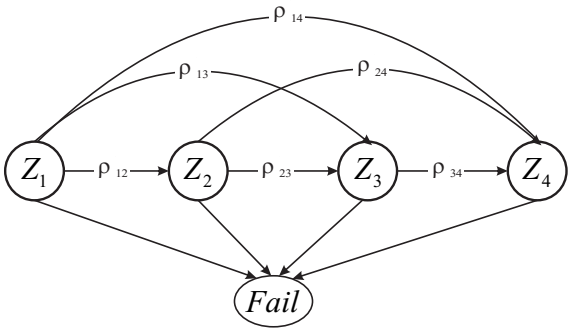

Figure 4.16: BN for the case of unequal correlation between components, for an example of four components.

For the case of uncorrelated components, the BN is illustrated in Figure 4.17, where there is no connection between component limit state functions.

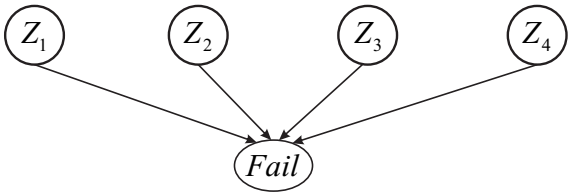

Figure 4.17: BN for the case of uncorrelated components, for an example of four components.

The BN can be sampled, and the system failure probability $P_{f,ss}$ estimated according to Eq. 4.32, but this evaluation would be as computationally intensive as what Sues and Cesare proposed ([52]). In the reliability framework developed in Chapter 2, the underlying random variables are explicitly included, which means efficient methods like MC Importance sampling can be applied to the load. The networks illustrated in Figs. 4.15-4.17 are not conducive to importance sampling because it would require replacing the distributions of all of the variables (Z_i) with biased distributions. Generally, the more corrections that need to be applied, the poorer the performance of importance sampling.

The system failure probability for the systems illustrated in Figs. 4.15-4.17 can also be calculated without sampling. The probability of failure of the system can also be expressed according to Eq. 4.33, which can be solved by evaluating the multivariate normal distribution.

$$P_{f,ss} = P\left(Fail = 1\right) = \frac{1}{N}\sum_{j=1}^{N} Fail_j \tag{4.32}$$

$$P_{f,ss} = P\left(\bigcup_{i=1}^{m} Z_i < 0\right) \tag{4.33}$$

The steps involved are adapted from Hohenbichler & Rackwitz ([58]), and provided in Eq. 4.34. The final step shows that once the linearization of Z has occurred (i.e. $Z = \alpha_1 U_1 + \alpha_2 U_2 + ...$), the series system failure probability can be calculated as $P_{f,ss} = 1 - \Phi_{\mathbf{m}}(\boldsymbol{\beta}; \Re_X)$, where Φ is the m-dimensional multivariate normal distribution, and $\Re_X$ is the correlation matrix of $\mathbf{X}$, which is equivalent to $\Re_Z$ as they differ only by a constant. At the time Hohenbichler & Rackwitz wrote their paper, and even at the time that the Equivalent Planes method was being developed, evaluating a large multivariate normal distribution took too much computing power to be feasible. However, more recent methods by Genz and Bretz ([70], [71]) have been adopted in statistical packages, such as Matlab (which was used in this research).

$$\begin{aligned}
P_{f,ss} &= P\left(\bigcup_{i=1}^{m} Z_i < 0\right) \\
&= P\left(\bigcup_{i=1}^{m} \alpha_i \mathbf{U} + \beta_i < 0\right) \\
&= P\left(\bigcup_{i=1}^{m} X_i < -\beta_i\right) \\
&= 1 - P\left(\bigcap_{i=1}^{m} X_i > -\beta_i\right) \\
&= 1 - P\left(\bigcap_{i=1}^{m} X_i < \beta_i\right) \\
&= 1 - \Phi_{\mathbf{m}}(\boldsymbol{\beta}; \Re_X)
\end{aligned} \tag{4.34}$$

I evaluated Eq. 4.34 using the multivariate normal distribution function in Matlab, which relies on the quasi-MC method described in [70], for the case of equicorrelated components (Case I). The results are shown in Figs 4.18-4.20 below. The quasi-MC method in Matlab for evaluating the multivariate normal performs well for reliability indices of 3 and 4, for all component correlations. For an index of 5, the method performs well until the correlation gets high (> 0.7), and for a reliabiity index of 6, the method only performs accurately when the components are barely correlated ($\rho = 0.2$). The poor performance appears to be due to an inability of the quasi-MC method to accurately calculate the error in the probability estimate in some cases. For example, for $\beta_c = 6$ and $\rho = 0.9$, and $m = 250$, the true system failure probability is about $5 \cdot 10^{-8}$, while the quasi-MC method returns an estimate of about $1 \cdot 10^{-9}$. However, the method estimates an error of about 10^{-12}, which is clearly incorrect. Details about the quasi-MC algorithm to estimate both system probabilities and error fall outside the scope of this dissertation. Based on the results of our application of it to Case I, I conclude that directly evaluating the multivariate normal distribution is still too unstable for practical application. I do

recommend keeping up with the literature, as new algorithms to evaluate multivariate normal distributions continue to be developed.

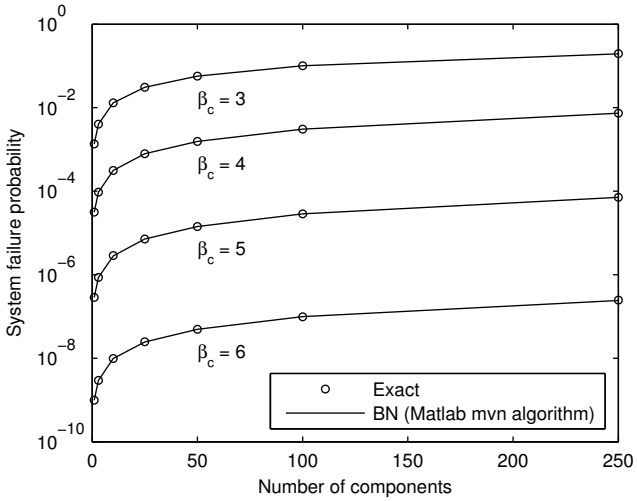

Figure 4.18: System failure probabilities using the multivariate normal function in Matlab, for a series system of equicorrelated components, with $\rho = 0.2$.

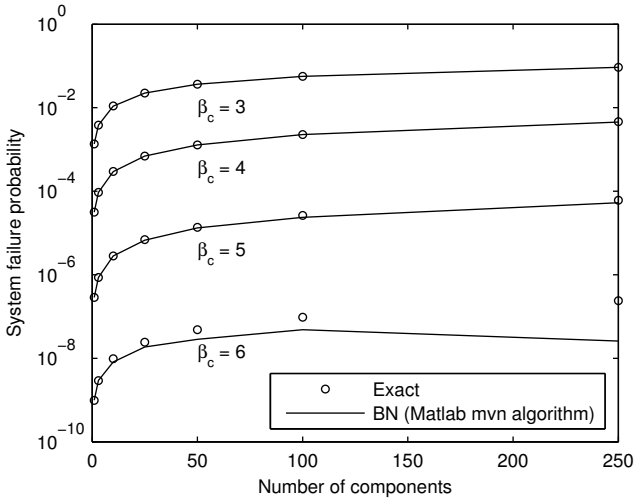

Figure 4.19: System failure probabilities using the multivariate normal function in Matlab, for a series system of equicorrelated components, with $\rho = 0.5$.

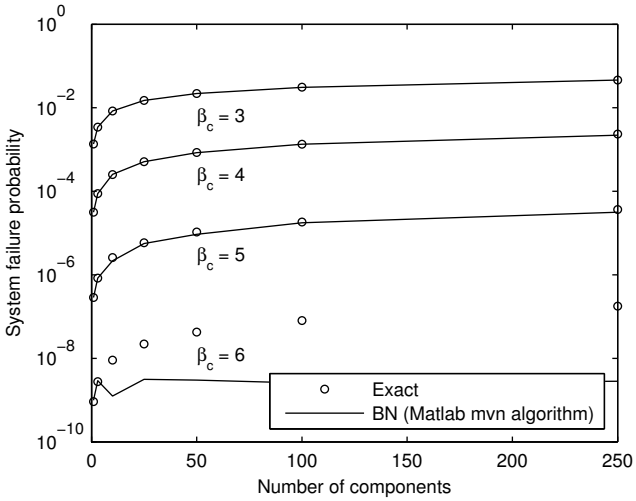

Figure 4.20: System failure probabilities using the multivariate normal function in Matlab, for a series system of equicorrelated components, with $\rho = 0.7$.

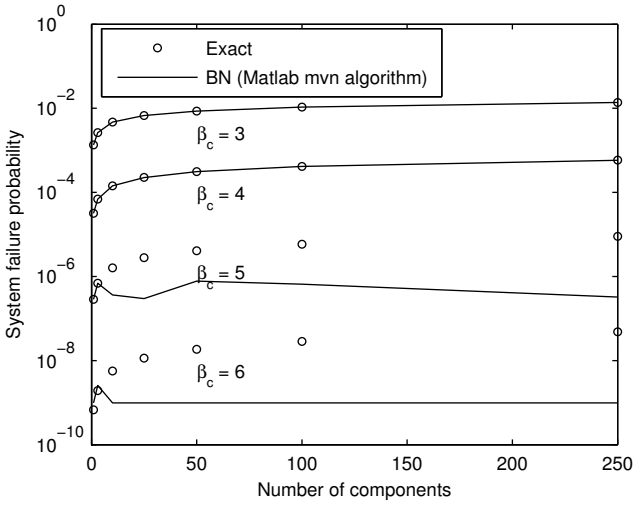

Figure 4.21: System failure probabilities using the multivariate normal function in Matlab, for a series system of equicorrelated components, with $\rho = 0.9$.

4.7. DISCUSSION AND CONCLUSIONS

The Equivalent Planes method has been used in Dutch system reliability modeling of flood defenses for decades. The reliability model is at the heart of national flood risk analysis, the results of which are used to drive major flood prevention policies in the Netherlands. The critical role of the model motivated this research to determine the

accuracy of the Equivalent Planes method, and under which situations unacceptable error may be encountered.

A Monte Carlo directional sampling method was used to compute 'exact' reliability estimates with which to compare the Equivalent Planes results. Table 4.11 shows the computation times[3] (in minutes) required for Monte Carlo directional sampling (with an imposed accuracy of 0.01 on the estimate of the system reliability index), for some of the system configurations that were tested as part of Case I (see section 4.4.2), for an inter-component correlation of 0.9. The Equivalent Planes method computed each of these configurations in less than 1 second. Consider a system with 250 components and component reliability indices of 5; in this case, the Monte Carlo directional sampling method required 1079 minutes – over 17 hours – compared with 0.95 seconds required by the Equivalent Planes method. In reality, systems are likely to have a large number of components, and component reliability indices that are greater than 5; thus, the reduction in computation time for real systems will be substantial. The accuracy of the EP method for the same configurations are shown in Table 4.12, given as a ratio of EP-calculated to exact failure probability. Tables 4.11 and 4.12 together indicate the trade-off between accuracy and computation time. Recall that for these configurations are for equi-reliable components; in real systems, where one or a few components have a lower reliability, the errors are expected to be less.

Table 4.11: Computation time (in minutes) for Monte Carlo directional sampling to compute the system reliability, for series systems of m = 3, 10, 50, 100, and 250 components with equal component reliability indexes (β_c) and equally correlated with a correlation coefficient of 0.9. In comparison, the EP method requires approximately 1 second for all cases shown in the table.

	$m=3$	$m=10$	$m=50$	$m=100$	$m=250$
$\beta_c=3$	<1	<1	1	1	3
$\beta_c=4$	<1	1	7	13	32
$\beta_c=5$	<1	8	188	424	1079

Table 4.12: Accuracy of the EP method, presented as the ratio of the EP-computed to exact system failure probability, for series systems of m = 3, 10, 50, 100, and 250 components with equal component reliability indexes (β_c) and equally correlated with a correlation coefficient of 0.9. The accuracy in this table can be compared with the computation times presented in Table 4.11.

	$m=3$	$m=10$	$m=50$	$m=100$	$m=250$
$\beta_c=3$	1.0	1.2	1.6	1.8	2.2
$\beta_c=4$	1.0	1.2	1.6	1.9	2.4
$\beta_c=5$	1.0	1.2	1.6	1.9	2.5

The error in the Equivalent Planes method for different system configurations was computed. When the components are not too correlated (e.g. a correlation coefficient up to about 0.5), the error in the method is generally negligible, particularly when the components have high reliability indices. Inaccuracies become apparent for large systems with highly correlated components, and for components with lower reliability in-

[3] Computation times are based on a 2.8 GHz computer with 8 GB RAM

dices. In all cases, the Equivalent Planes system failure probability estimates were within a factor of three times the correct system failure probability. Recall that these results are for extreme system configurations in which the components all have equal reliability indices. In reality, a few components will likely dominate the failure probability, and the error will be much lower. Furthermore, even three times the correct failure probability can be quite negligible for systems with very small failure probabilities. For example, consider a system of 250 equi-correlated components, with component reliability indices of 6, correlated with a coefficient of 0.9; the true failure probability is 4.81E-8 and the estimate is 1.14E-7, for a factor difference of 2.4 (see Table 4.3). In many applications, where the probability needs to be below a certain safety standard, this difference will not be important. Furthermore, other uncertainties in the reliability analysis – for example, due to the parameterization of the random variables contributing to failure – will likely overshadow this small error, making it essentially negligible. Tolerable error was discussed, and an example was provided for a levee system with loads similar to those found in the river regions of the Netherlands. It is important that researchers investigate tolerable error for their specific case to determine if the Equivalent Planes method will be sufficiently accurate. When it is, it is a very attractive method, particularly when considering the gain in computational time over more exact methods.

5

BAYESIAN UPDATING OF PIPING FAILURE PROBABILITIES

Probability then, refers to and implies belief, more or less, and belief is but another name for imperfect knowledge, or it may be, expresses the mind in a state of imperfect knowledge.

A. DeMorgan

5.1. INTRODUCTION

Estimates of levee system reliability can conflict with experience and intuition. For example, a very high failure probability may be computed while no evidence of failure has been observed, or a very low failure probability when signs of failure have been detected. This conflict results in skepticism about the computed failure probabilities and an (understandable) unwillingness to make important management decisions based upon them. However, although intuition is an important guide, it can also be misleading. Not all observations are informative, and it is important to have a quantitative tool to estimate the change in failure probability that results from a survival observation. BNs are useful in these circumstances because they allow us to use observations to improve our reliability estimates quantitatively.

In this chapter the BN methodology described in Chapter 2 is applied to two systems of levees in the Netherlands, a regional system and a primary one. In the Netherlands, regional levee systems offer protection from smaller (often regulated) waterways, such as canals. Primary levee systems protect against large bodies of water, like rivers, lakes, and the sea. Safety standards in the Netherlands are expressed in terms of a failure probability that must not be exceeded. For the regional levees, this failure probability is around 1/100. The primary levees have much more stringent standards, because their breach

would lead to catastrophic consequences. The failure probabilities of these levees must be lower than 1/1250 in the river regions to 1/10000 along the coast. Both case studies consider failure probabilities of the levee system due to the mechanism piping, in which pressure differences between the waterside and land side of the levee can cause a pipe to form under the levee. In both cases, survival observations are used to update the system failure probability.

One of the aims of this dissertation is to verify key algorithms in Hydra-Ring, and compare its efficiency to that of the BN methodology. To this end, the Hydra-Ring algorithms are also applied to the primary levee system, to calculate both segment and system failure probabilities which can be compared to the BN estimates. For the Hydra-Ring approach, the modified outcrossing (MO) algorithm is used to calculate the segment failure probability (by accounting for the length effect; see Chapter 3), and the Equivalent planes (EP) method is used to calculate the system failure probability by combining the correlated segment failure probabilities (see Chapter 4 for details about the EP method). Differences in both segment and system failure probabilities, as well as computation times, are investigated and discussed.

A key question in reliability updating is: When are observations useful? Not all observations will have equal impact on the failure probability, and it is important to know in advance what factors will influence the impact, and if the case at hand is a good candidate for updating. Implementation of either method - the BN or the Hydra-Ring algorithms - takes time and care, and one would prefer to dedicate such time only in cases when the reduction in failure probability is non-negligible. This chapter therefore addresses the question: under which conditions are survival observations useful?

The structure of the chapter is as follows. The piping failure mechanism is described in Section 5.2. The BN for failure of a levee system due to the piping failure mechanism is described and presented visually in Section 5.3. Section 5.4 presents the application of the BN to the regional levee system, and Section 5.5 presents the application of the BN and the MO and EP algorithms to the primary levee system. Section 5.7 discusses which survival observations are useful in general, and summarizing conclusions are provided in Section 5.8.

5.2. PIPING FAILURE MECHANISM

In the two case studies presented in this chapter, I consider failure probabilities due to the piping failure mechanism. Figure 5.1 provides an illustration that supports the following description. When the pressure difference between the outside water level (h in Figure 5.1) and the landside water level (h_{ls}) is great enough, it can increase the soil pore water pressure in the aquifer (sand layer) to the point that it causes the clay layer to uplift (i.e. rupture) on the landside of the levee. Once this occurs, if the pressure difference is great enough, sand can begin to transport from the aquifer onto the landside of the levee. What follows is an eroded pipe within the aquifer, allowing water from the landside of the levee to start filling in the pipe, as sand continues to erode. If the pipe reaches the waterside, the levee will essentially be resting on a film of water, which is a very unstable situation, and is likely to lead to collapse of the levee.

The piping mechanism considered in this application is described by two limit state functions, one describing uplift of the clay layer (uplift) and the other describing the

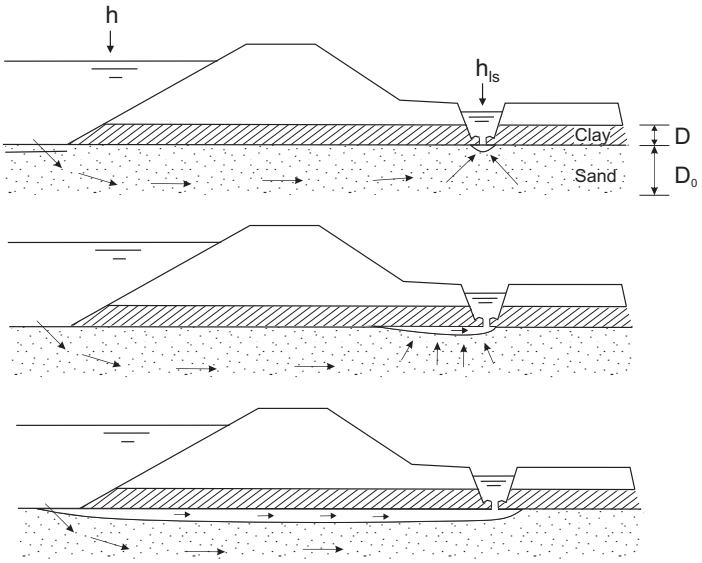

Figure 5.1: Progression of the piping mechanism, beginning with uplift (top) until the pipe is complete (bottom)

initiation of the pipe formation (piping). Failure is considered to occur if both the uplift and piping limit state function are negative (parallel system).

UPLIFT

The limit state function for uplift (Z_u) is given in Equation 5.1. This formula is widely used in the Netherlands and was recommended by the Technical Advisory Committee on Flood Defenses ([72]). It compares the actual water level difference, $h-h_{ls}$, with a critical water level difference, h_c. The variable m_u is a model factor, accounting for uncertainty in the formula for h_c (see Equation 5.2). The factor m_h is known as a damping factor, and accounts for potential loss of pressure between the water levels on the water-side and land-side of the levee.

$$Z_u = m_u h_c - m_h (h - h_{ls}) \tag{5.1}$$

The critical water level difference, h_c, depends on the thickness of the impervious (clay) layer (D), the volumetric weight of wet clay (γ_{wc}), and the volumetric weight of water (γ_w).

$$h_c = \frac{\gamma_{wc} - \gamma_w}{\gamma_w} \cdot D \tag{5.2}$$

PIPING

The limit state function for piping, Z_p, was developed ([73]) and recently improved ([74]) by Sellmeijer. The limit state function, and supporting equations, are provided in Eqs. 5.3 through 5.6. All variables are defined in Table 5.1.

$$Z_p = m_s h_p - (h - 0.3D - h_{ls}) \tag{5.3}$$

The critical head difference h_p is given in Equation 5.4). The two coefficients, α and c, are provided in Equation 5.5 and 5.6, respectively. For some physical intuition: α reflects the finite thickness of the sand layer and c has to do with erosion characteristics of the sand.

$$h_p = \alpha \cdot c \cdot L \left(\frac{\gamma_k - \gamma_w}{\gamma_w} \right) \cdot \tan(\theta) \tag{5.4}$$

$$\alpha = 0.91 \left(\frac{D_0}{L} \right)^{\frac{0.28}{\left(\left(\frac{D_0}{L} \right)^{2.8} - 1 \right)} + 0.04} \tag{5.5}$$

$$c = \eta \cdot d_{70_m} \cdot \left(\frac{g}{\nu \cdot k \cdot L} \right)^{\frac{1}{3}} \cdot \left(\frac{d_{70}}{d_{70_m}} \right)^{0.4} \tag{5.6}$$

5

Table 5.1: Description and distribution types of the variables used in the piping analysis; logn = lognormal, norm = normal, det = deterministic

Variable	Description	Distribution	Constant over segment
D_0	Thickness of aquifer	logn	no
D	Thickness of blanket layer	logn	no
L	Distance, waterside levee toe to landside water	logn	no
θ	Bedding angle of sand	norm	no
d_{70}	70th-percentile of sand grain diameter	logn	no
η	Drag coefficient	logn	yes
γ_{wc}	Volumetric weight of blanket layer	logn	no
γ_k	Volumentric weight of sand	logn	no
m_u	Error in critical pressure difference, for uplift	logn	yes
m_h	Damping factor	logn	yes
m_s	Error in piping model (Sellmeijer)	logn	yes
k	Permeability of aquifer	logn	no
h_{ls}	Water level on landside of levee	norm	yes
d_{70m}	Reference value for d_{70}	det	–
g	Gravitational constant	det	–
γ_w	Volumetric weight of water	det	–
ν	Viscosity of water	det	–

5.2.1. SURVIVAL OBSERVATIONS FOR THE PIPING MECHANISM

This section describes explicitly what is meant here by a survival observation, as it relates to the piping failure mechanism. Uplift refers to the rupture of the blanket layer due to upward pressure in the aquifer. This phenomenon is easy to confirm; if a rupture is seen, or substantial seepage is observed, it can be concluded that uplift has occurred and that the limit state Z_u (see Eq. 5.1) must be negative. Or conversely, when no rupture or seepage is observed, Z_u must be positive. When no pipe has reached completion (i.e. become observable), and the levee shows no signs of weakness, it is assumed in this chapter that the limit state function for piping Z_p (see Eq. 5.3) is positive. Therein lie two somewhat unconservative assumptions. The first is that the duration of the observed high load was sufficient for a pipe to reach completion. If it was not, then the same load may result in failure for a longer duration, but this is not accounted for in the methodology. The

second is that (unseen) damage from a previous load does not contribute to the rate of progression of a pipe during future high loads. In reality, prior (partial) internal erosion will result in the completion of a pipe forming at a lower load than if that prior erosion had not taken place. For the purpose of illustrating the application and abilities of a BN, assuming $Z_p > 0$ for the observation of 'no failure' is useful. However, for practical application, it should be discussed if a more conservative approach, such as that taken in [50], is more appropriate. In that case, Z_p was not assumed positive when survival was observed. Rather, only the uplift limit state was assumed known under survival conditions. The impact of updating is in that case much less, but is more conservative.

5.3. BN FOR PIPING

The structure of the BN for piping is dictated by the formulaic representation of piping, as described in Equations 5.3 through 5.6. The variables that play a role in the piping mechanism, which are described in Table 5.1, are the input random variable nodes in the BN. Table 5.1 also indicates whether a variable is constant over the length of the segment. If so, it will be represented by one node per segment in the system. The variables that are not constant are spatially variable and will be represented by n nodes, where n is the number of cross sections representing the segment. Figure 5.2 shows the BN for a single cross section. No spatial variability occurs within a single cross section, but I have illustrated which variables will be subject to spatial variability when the number of cross sections increases. These are the variables at the top of the network, with the superscript 1 to indicate they are the variables in the first cross section. The limit state functions for piping and uplift, Z_p and Z_u, are connected to a failure node *Fail*, which is 1 when $Z_p < 0 \cap Z_u < 0$, and 0 otherwise.

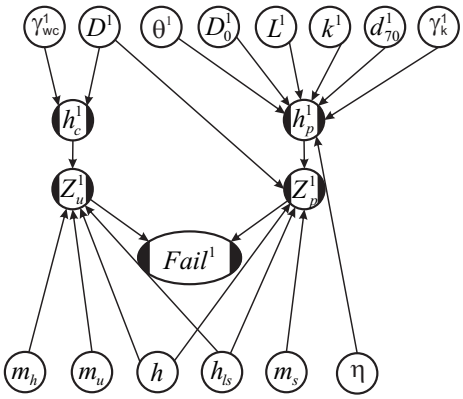

Figure 5.2: BN for a single cross section, for the failure mechanism piping. Clear circular nodes represent input random variables, and nodes with black edges represent functional variables (see Eqs. 5.3 through 5.6).

The BN for an example segment represented by three cross sections is illustrated in Figure 5.3. The spatially-variable random variables are connected via arcs to the same variables in the other cross sections. For example D^1 is connected to D^2 and D^3. The variables which are constant over the segment are again displayed at the bottom of the

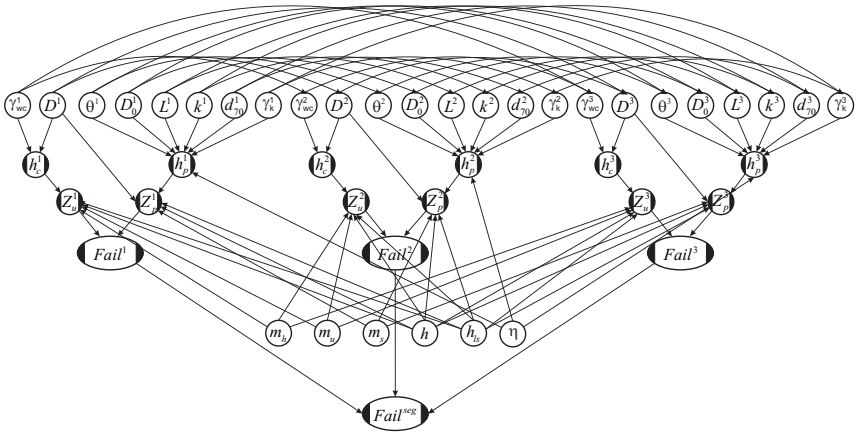

Figure 5.3: BN for a segment represented by three cross sections, for the failure mechanism piping. Clear circular nodes represent input random variables, and nodes with black edges represent functional variables (see Eqs. 5.3 through 5.6).

network, and are only represented one time per segment.

Arcs in the network that lead into functional nodes are described by the formulas in Eqs. 5.3 through 5.6. Arcs between input random variables (such as D^1 and D^2) are specified by a product moment correlation coefficient. These are calculated using the positive definite correlation function given in Eq. 5.7, where Δx_{ij} is the distance between a variable in the i^{th} and j^{th} cross section, and d_x is a parameter that dictates how quickly the spatial correlation decreases to zero. The values for d_x are provided in each of the case studies.

$$\rho_{ij} = \exp\left(\frac{\Delta x_{ij}}{d_x}\right)^2 \tag{5.7}$$

The failure nodes for each cross section are connected to the failure node for the segment, $Fail^{seg}$, which is 1 whenever $Fail^i = 1$, for any $i = 1, ..., n$. The system representation is similar to that of the segments, with a couple of caveats. First, I make an assumption that spatial correlation of the soil variables between segments is zero (i.e. the variables D^i in segment 1 are not connected to the D^i in segment 2, for any i). This is a reasonable assumption because segments are typically delineated due to clear changes in soil type/stratification. Second, of the variables that are constant over the segment, a few of these may also be constant over the system. In our case studies, the canal or river water level is assumed constant over the length of the system, as well as the model uncertainty parameters m_u and m_s and White's constant η. These system constant variables are represented by a single node for the whole system, whereas m_h and h_{ls}, which are constant over the segment but not over the system, are represented n_{seg} times, where n_{seg} is the number of segments in the system.

5.4. REGIONAL LEVEES

A map of the levee system protecting Heerhugowaard, including the levee segment numbers, is presented in Figure 5.4. It is split into 18 levee segments. The levee system considered in this application is composed of just three of these segments (9, 11, and 12). The reason for this is that only the southern and western levee segments (segment 5 through 12 in Figure 5.4) cause significant damage if they fail, and of these segments, 9, 11, and 12 were the real 'weak links' in the sense that they had substantially higher computed failure probabilities than the others. The water board responsible for the levees is highly skeptical about these failure probabilities, because they have never observed any evidence of piping. This made it an ideal case to apply a Bayesian network and make use of its capabilities to incorporate a survival observation.

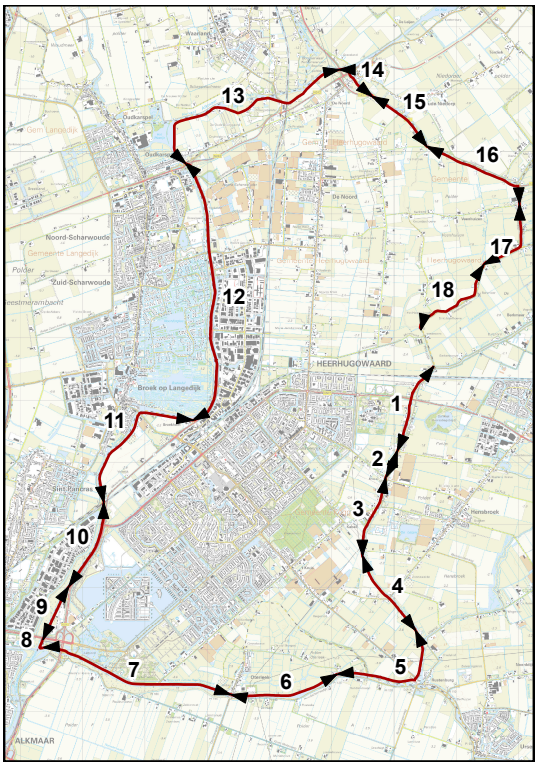

Figure 5.4: Location of the levee system around Heerhugowaard, with 18 segments

5.4.1. DATA

This section describes the prior probability distributions of the variables in the Bayesian network. Table 5.2 shows the distribution parameters for all of the variables relevant for the piping limit state functions (with the exception of the canal water level).

The water level in the canal is regulated; when needed, water is pumped into the canal from the lower-lying protected area, as long as the water level in the canal does

not exceed the maximum tolerated level. In this canal, that level is exactly equal to the Dutch datum, known as Amsterdam Ordinance Datum (AOD).

A two-parameter generalized Pareto distribution (GPD) was fitted to independent water level peaks above a selected threshold. The threshold was chosen according to the method described in [75]. The parameters of the GPD are the shape (σ) and scale (ξ) parameters; the distribution function is presented in Eq. 5.8. After fitting the GPD to the data, the distribution was then modified so that any water level above the maximum tolerated level had an exceedance probability of zero. In this way, the regulated aspect of the canal is taken into account. It is possible that human error may lead to a water level above the maximum tolerated level, but that was not considered in this case study. Figure 5.5 shows the exceedance probability curve for the water level, and Table 5.3 shows parameters of the GPD.

$$F(x) = \begin{cases} 1 - \left(1 - \xi \dfrac{x}{\sigma}\right)^{1/\xi}, & \xi \neq 0 \\ 1 - \exp\left(-\dfrac{x}{\sigma}\right), & \xi = 0 \end{cases} \tag{5.8}$$

5

Table 5.2: Input values for the three segments (S9, S11, and S12). Shown are the mean M of the distributions, the standard deviation SD, and the correlation length d_x (in meters). Note that for d_x, ∞ means the variable is fully correlated over the segment.

Variable	Units	M (S9)	M(S11)	M (S12)	SD	d_x
D_0	[m]		15		$0.1M$	200
D	[m]	0.3	0.01	0.01	$0.1M$	200
L	[m]	37.5	17.75	39	$0.1M$	3000
θ	[deg]		37		3	600
d_{70}	[m]	3.15E-04	2.62E-04	3.15E-04	$0.15M$	180
η	[-]		0.25		$0.05M$	∞
γ_{wc}	[kN/m^3]		16		$0.05M$	300
γ_k	[kN/m^3]		18		$0.05M$	300
m_u	[-]		1		$0.1M$	∞
m_s	[-]		1		$0.1M$	∞
m_h	[-]		1		$0.08M$	∞
k	[m/s]	1.74E-04	9.26E-05	1.74E-04	M	600
h_{ls}	[m+AOD]	-3.6	-3.6	-2.85	$0.1M$	∞
d_{70m}	[m]		2.08E-04		–	–
g	[m/s^2]		9.81		–	–
γ_w	[kN/m^3]		10		–	–
ν	[m^2/s]		1.00E-05		–	–

Table 5.3: Shape and scale parameters of the Generalized Pareto distribution (GPD) of canal water levels, the threshold above which canal water level peaks were selected, and the number of peaks used to fit the GPD parameters.

Shape	Scale	Threshold	# Peaks
0.0651	0.0318	-0.3984	100

5.4.2. PRIOR ANALYSIS WITH THE BN: REGIONAL CASE

For each segment in the system, the number of cross sections in the segment BN was iteratively increased until the failure probability of the segment reached an asymptote.

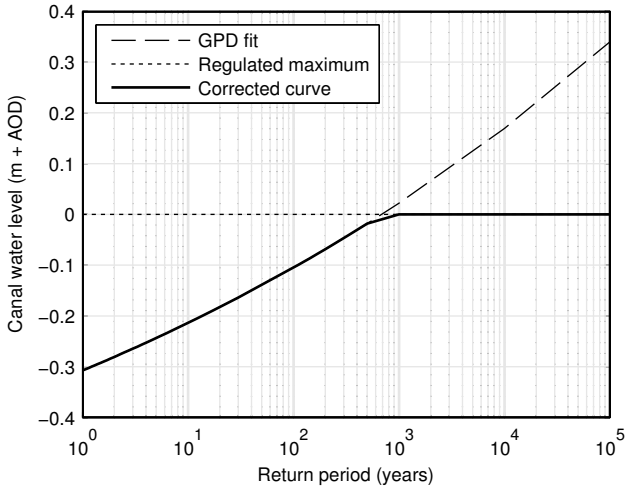

Figure 5.5: Water level exceedance probabilities for the water level in the Schermer Canal

The stop criterion was defined such that the iterative procedure stops when the segment failure probability of the previous 10 iterations all lie within the 95% confidence interval of the current iteration. Figs. 5.6 - 5.7 show the prior failure probability of the three segments under consideration, as a function of the number of cross sections representing the segment. The 95% confidence intervals around the final estimate is also shown. Table 5.4 summarizes the results, providing the cross section failure probability, the selected number of cross sections to represent the segment, and the segment failure probability for each of the three segments.

Table 5.4: Prior cross sectional failure probability ($P_{f,CS}$), prior segment failure probability ($P_{f,Seg}$), number of cross sections (# CS), and the length of the segment (L_{seg})

Segment	$P_{f,CS}$	$P_{f,Seg}$	# CS	L_{seg} (m)
9	0.32	0.59	47	1000
11	0.85	0.99	47	2000
12	0.12	0.57	75	4000

After determining the number of cross sections that will represent each of the segments, the levee system network can be built. To do this, variables which are fully correlated over the system need to be identified. Table 5.2 provided the correlation lengths (d_x) for each of the variables; those with a value of infinity (∞) are fully correlated over the entire segment. However, only a few of these are also fully correlated over our entire system. These were the water level in the canal (h), the model error in the uplift and piping models (m_u and m_s, respectively), and the drag coefficient (η). All other variables were taken to be uncorrelated between segments.

The BN estimated a system failure probability of 0.998, which is unsurprising, given the failure probability of 0.99 of Segment 11 (see Table 5.4). This probability is highly

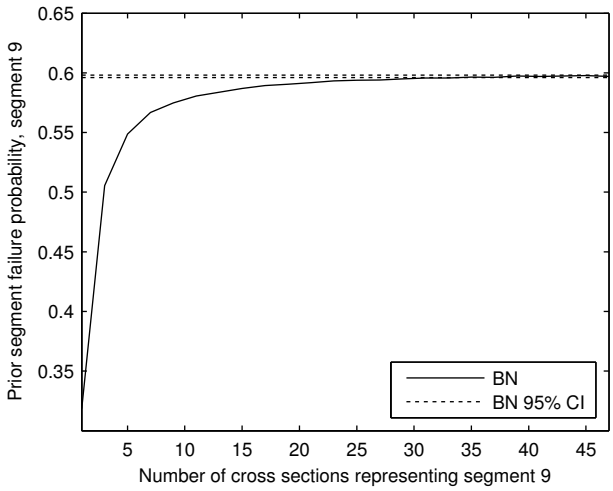

Figure 5.6: Prior segment failure probability as a function of number of cross sections representing the segment, shown here for levee segment 9

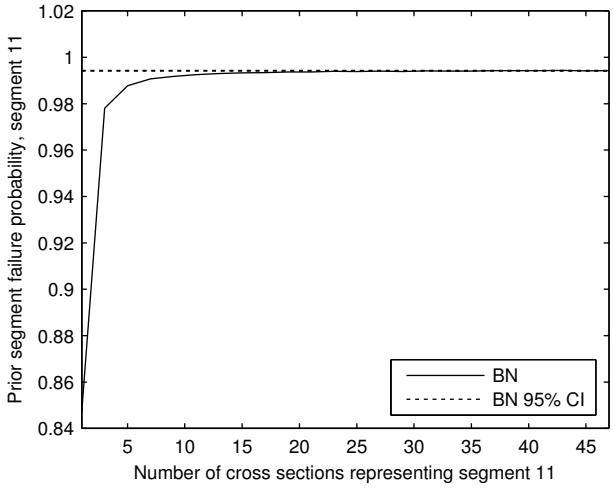

Figure 5.7: Prior segment failure probability as a function of number of cross sections representing the segment, shown here for levee segment 11

suspect, given that no evidence of piping has been observed in tens of years. This made it a good candidate for updating using a survival observation.

5.4.3. INCORPORATING A SURVIVAL OBSERVATION: REGIONAL CASE

The highest water level on record for this stretch of levee is -0.1460 AOD, which corresponds to a return period of about 40 years, and was observed in January 2003. The

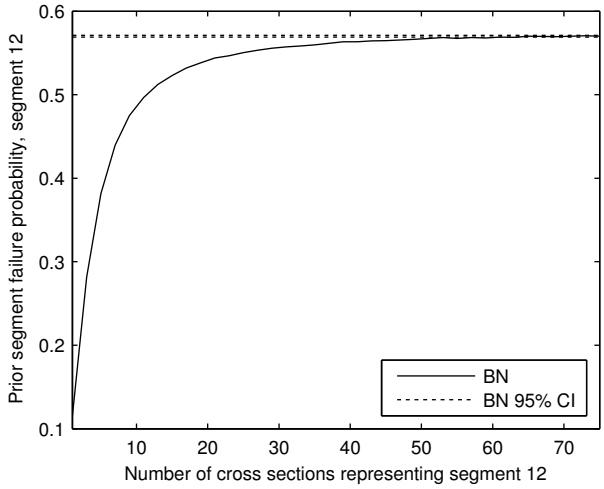

Figure 5.8: Prior segment failure probability as a function of number of cross sections representing the segment, shown here for levee segment 12

levee survived with no evidence of weakness. The coupled survival observation is therefore $h = -0.146$ and $Z_{p,i} > 0 \cap Z_{u,i} > 0$, $i = 1, ..., n$, for all n cross sections in the segment (or system when doing a posterior analysis at the system level). The posterior segment failure probabilities were calculated for each of the segments separately, and for the system as a whole.

As described in Chapter 2, inference is performed using rejection sampling. Specifically, h is set to a deterministic value h_{obs}. The network is then sampled, and only those samples for which $Z_{p,i} > 0 \cap Z_{u,i} > 0$, for $i = 1, ..., n$. The retained samples form the posterior joint distribution, which are used to calculate the updated failure probability. Details about how this is done are provided in Chapter 2 (see Section 2.3.5). The posterior segment failure probabilities are presented in Figs. 5.9-5.11. Note that for segment 11 the 95% confidence interval is extremely wide. This is because the number of posterior samples is very low. Recall that this segment had a prior failure probability of 0.99 (see Table 5.4). This means that with rejection sampling, only 1% of the samples are retained for posterior analysis. I started with $1E6$ samples, so the posterior failure probability is computed with only about $1E3$ samples.

The BN for the system returned a prior failure probability of 0.998. After incorporating the survival observation, the system failure probability estimate is 0.0045. This is a reduction of over two orders of magnitude.

The posterior analysis introduces dependence between previously independent soil variables. Therefore, if the updated distributions are used in other analyses (e.g. another failure mechanism that depends on some of the same variables), this posterior dependence must be taken into account. The posterior marginal distributions of the variables that were most affected by the survival observation are presented in Figure 5.12 for Segment 11. The other segments showed similar results, although this segment had the

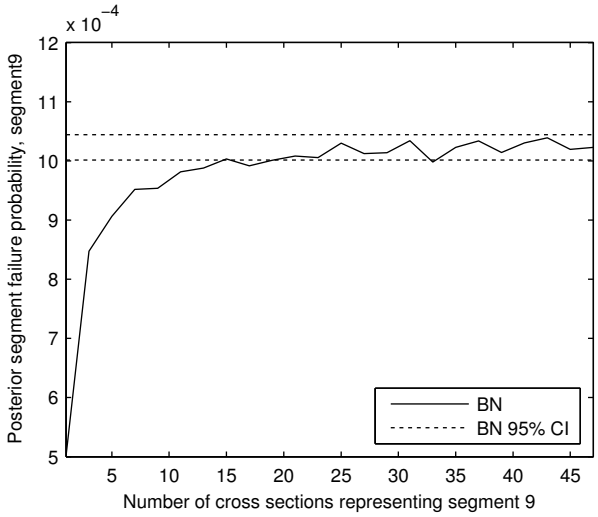

Figure 5.9: Posterior failure probability for segment 9

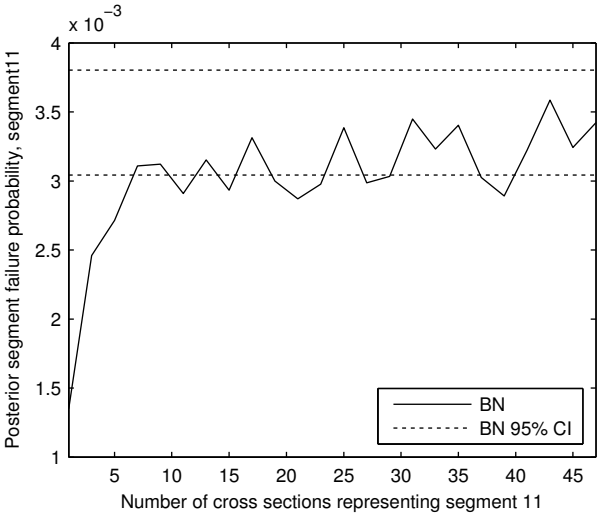

Figure 5.10: Posterior failure probability for segment 11

most impact as its prior failure probability was so high.

5.4.4. DISCUSSION

Discussions with the water board following the completion of this research illuminated an interesting aspect missing from the Bayesian network representation of the system. The piping mechanism depicted in Figure 5.1 assumes that a sand layer underlies the

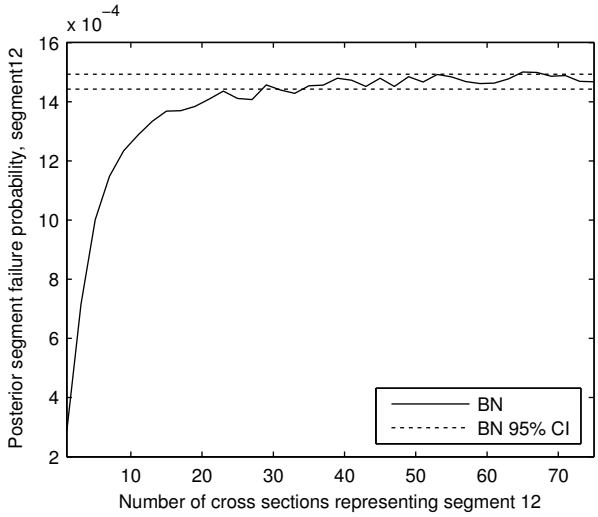

Figure 5.11: Posterior failure probability for segment 12

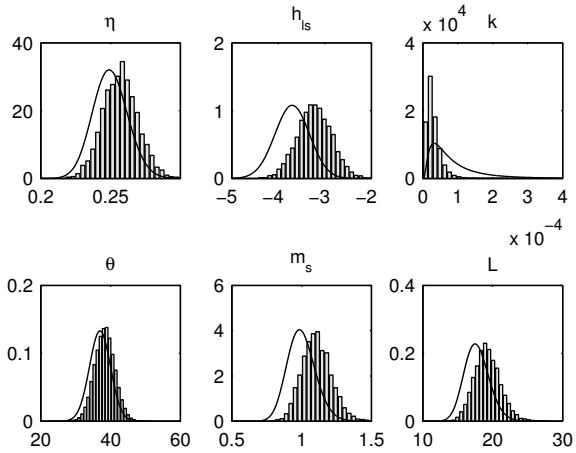

Figure 5.12: Updated distributions of a selection of soil variables

canal. In practice, there may be a clay layer between the canal and the aquifer, essentially making the soil water pressure in the aquifer immune to the water level in the canal, and rendering the piping mechanism impossible. Field measurements for the Heerhugowaard system have concluded that this clay layer exists at a number of locations. This can be accounted for by including a node that represents the existence (or nonexistence) of such a clay layer. The node would act like a switch; when the layer exists, the failure probability of the segment will be zero, regardless of the value of the limit state functions for piping. When the layer does not exist, the limit state functions

will determine whether failure occurs. Inclusion of the node was preliminarily explored and results show that it has a dominant influence, with very high posterior probabilities of existence. It also has the effect of reducing the posterior failure probability relative to the case when the clay layer is not accounted for. Therefore, inclusion of this node should be done in close conversation with those responsible for the safety of the region. One of the issues that specifically needs to be addressed is whether it is acceptable to depend on the mucky clay layer, which is not a managed part of the defense system, to ensure the safety of the inhabitants.

5.5. PRIMARY LEVEES

The primary system considered contains three levee segments, 13a, 13b, and 14, along the IJssel River in the Netherlands. The system protects agricultural land, as well as the part of the picturesque city of Zutphen that lies to the west of the IJssel River. Figure 5.13 presents the location of the system. The failure probability of the middle segment - segment 13b - was estimated to be too high in a national flood risk analysis study ([47]). Those familiar with the system were skeptical about this high failure probability, which is why it was selected for a case study. The surrounding segments, 13 and 14, had lower failure probabilities, but were included in this study to illustrate the application to a system.

5.5.1. DATA

The soil data used for the three segments was the same data used for the national flood risk analysis, carried out with the reliability model PC-Ring. The data is provided in Table 5.5. The mean and standard deviation of each of the variables is given, as well as the correlation length d_x, which dictates how quickly the spatial autocorrelation decreases. The distribution types were provided earlier in this chapter in Table 5.1.

Table 5.5: Input values for the three segments of the primary levee system (13a, 13b, and 14). Shown are the mean M of the distributions, the standard deviation SD, and the correlation length (in meters) d_x. Note that for d_x, ∞ means the variable is fully correlated over the segment.

Variables	Units	M (13a)	M(13b)	M (14)	SD(13a)	SD(13b)	SD(14)	d_x
D_0	[m]	9.4	10.2	8.7		0.3M		200
D	[m]	1.6	0	1.8	0.2M	–	0.08M	200
L	[m]	44.4	52.0	37.6	0.03M	–	0.09M	3000
θ	[deg]		43			3		600
d_{70}	[m]		4.0E-04			0.27M		180
η	[-]		0.3			0.1M		∞
γ_{wc}	[kN/m^3]	19.3	19.47	19.16		0.1M		300
γ_k	[kN/m^3]		27			0.01M		300
m_u	[-]		1			0.1M		∞
m_s	[-]		1			0.0.08M		∞
m_h	[-]		0.8			0.1M		∞
k	[m/s]	5.8E-04	5.4E-04	5.84E-04		1.3M		600
h_{ls}	[m+AOD]	7.0	6.2	6.5		0.1M		∞
d_{70m}	[m]		2.08E-04			–		–
g	[m/s^2]		9.81			–		–
γ_w	[kN/m^3]		10			–		–
ν	[m^2/s]		1.00E-05			–		–

The water level distribution was fitted using three quantiles that had been calculated

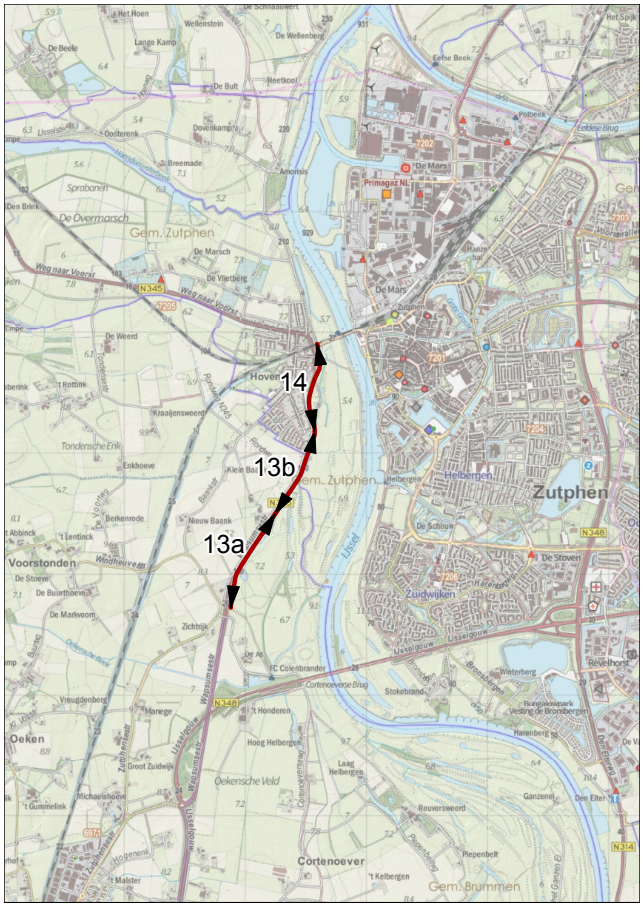

Figure 5.13: Location of the primary levee system

with PC-Ring. In the PC-Ring model, upstream discharges at Lobith (on the border with Germany) are sampled, and then used as an upstream boundary in a hydraulic model that calculates the associated water levels in the river branches. Because that has already been done, the hydraulic model was not needed in this analysis. Instead, three quantiles of the local water level were used which PC-Ring had previously calculated. These quantiles were the 1/125, 1/1250, and 1/12500 water levels. A Gumbel extreme value distribution (for maxima) was assumed, the formula for which is provided in Eq. 5.9. It has two parameters, a location parameter μ and a scale parameter β, which were fitted using the known quantiles. The parameters of the Gumel distribution, as well as the mean and standard deviation are provided in Table 5.6.

$$F(x) = \exp\left[-\exp\left[-\frac{1}{\beta}(x-\mu)\right]\right] \tag{5.9}$$

Table 5.6: Gumbel distribution parameters, IJssel river at location of segements 13a, 13b, and 14.

Location (μ)	Scale (β)	Mean	Standard deviation
7.36	0.31	7.5	0.40

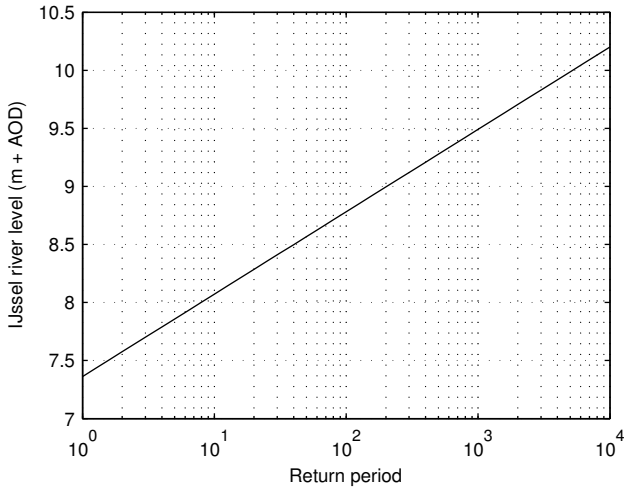

Figure 5.14: Distribution of the water level in the IJssel River at the location of the levee system

5.5.2. PRIOR ANALYSIS: PRIMARY CASE

As with the regional system, the number of cross sections in each segment BN was iteratively increased until the failure probability of the segment reached an asymptote. The same stop criterion was used, where the iterative procedure stops when the segment failure probability of the previous 10 iterations all lie within the 95% confidence interval of the current iteration. In this application, importance sampling was used to sample the river water level in the prior analysis (see Chapter 2 for details of importance sampling). This was not necessary in the regional case, because the prior failure probabilities were so high. As a biased distribution, the Gumbel distribution was used with the same scale parameter (see Table 5.6), but with a location parameter increased by 1 for segments 13a and 14, and by 0.5 for segment 13b. Table 5.7 summarizes the results, providing the cross section failure probability, the selected number of cross sections to represent the segment, and the segment failure probability for each of the three segments.

Table 5.7: Cross sectional failure probability ($P_{f,CS}$), segment failure probability ($P_{f,Seg}$), number of cross sections (# CS), and the length of the segment (L_{seg})

Segment	$P_{f,CS}$	$P_{f,Seg}$	# CS	L_{seg} (m)
13a	1.9E-5	6.8E-5	19	875
13b	3.9E-4	1.4E-3	43	763
14	1.7E-4	5.8E-4	35	770

PRIOR SEGMENT FAILURE PROBABILITY: COMPARISON OF BN AND MO METHODS

In Chapter 3, I compared the BN with the modified outcrossing (MO) method estimates of the segment failure probability, for a few numerical examples. In this application, the segment failure probabilities were also calculated using the MO method, and the results are presented together with the BN results. Figs. 5.15 - 5.17 show the prior failure probability of the three segments, as a function of the number of cross sections representing the segment. The 95% confidence intervals around the BN estimate is also shown. The MO estimate is shown as a horizontal line because it is not dependent on the number of cross sections representing the BN. The agreement between the two methods is not as strong as it was for the numerical example in Chapter 3, but it is still quite good. In terms of reliability index, the differences are only 2%, 1%, and 3% for segments 13a, 13b, and 14, respectively. Tabulated results of the estimates are presented together with the posterior results in Table 5.8.

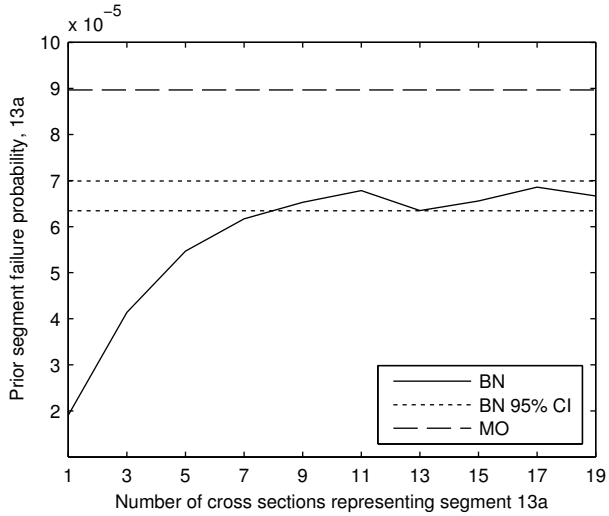

Figure 5.15: Prior segment failure probability for segment 13a, computed with the BN and the MO method. The 95% confidence interval around the BN estimate is also presented.

PRIOR SYSTEM FAILURE PROBABILITY: COMPARISON WITH EQUIVALENT PLANES METHOD

I compared the system failure probability calculated by the BN with the estimate derived using the Equivalent Planes (EP) method described in Chapter 4, which combines correlated system components (in this case our three 'components' are the three segments in our system). To use the EP method, the influence coefficients of the random variables are required. The cross-sectional influence coefficients are easily obtained when using FORM to calculate the failure probability, but the coefficients for the variables at the segment scale are not guaranteed to be the same. In fact, they will almost definitely *not* be the same because the more spatially variability a variable exhibits, the more likely it will be to have a weak realization and cause failure. Therefore, spatial variability will play a role in which variables are most influential at the segment scale. Though not

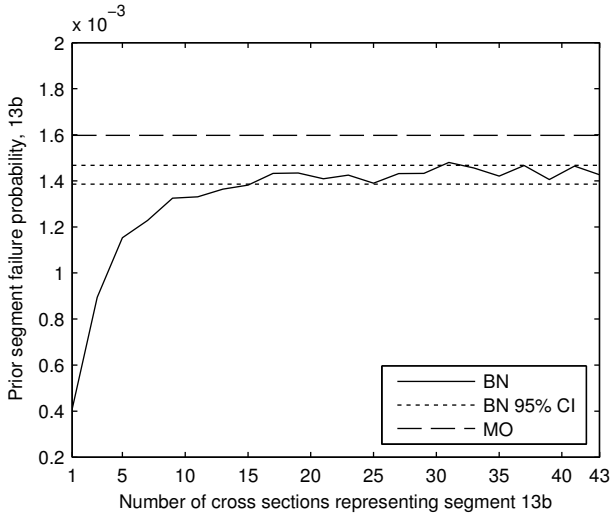

Figure 5.16: Prior failure probability for segment 13b, computed with the BN and the MO method. The 95% confidence interval around the BN estimate is also presented.

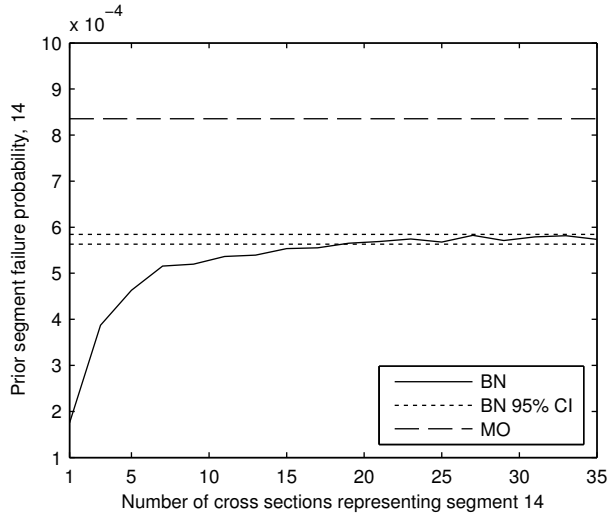

Figure 5.17: Prior failure probability for segment 14, computed with the BN and the MO method. The 95% confidence interval around the BN estimate is also presented.

described in this dissertation, the MO method implemented in Hydra-Ring is coupled with a method to calculate the influence coefficients at the segment scale. For details about how this is done, the reader is referred to the Hydra-Ring technical document ([11]). Thereafter, the EP method is applied as described in Chapter 4. The prior system failure probability estimate calculated with the EP method is referred to as a com-

bined MO/EP estimate, since the segment failure probabilities were calculated with the MO method and the system probability with EP. The BN estimate for the system failure probability was 1.9E-3; the MO/EP estimate was 2.5E-3. This is about a 30% difference, which is considered relatively minor at this magnitude. In terms of reliability index, the estimates are 2.9 and 2.8 for the BN and MO/EP methods, respectively, which equates to a difference of about 3%.

5.5.3. INCORPORATING A SURVIVAL OBSERVATION: PRIMARY CASE

For the survival observation, I considered two water levels: one with a return period of 400 years and one with a return period of 40 years. The first one is a very extreme water level of 9.2 m + AOD. The second is more realistic to have been observed, at 8.5 m + AOD. Both water levels are considered to get a feel for how the extremity of the observed load will influence the impact that a survival observation has on the failure probability. Figures are only presented for the case that h_{obs} = 9.2 m + AOD, which is the 1/400 year water level, but tabulated results are presented for both observed water levels.

POSTERIOR SEGMENT FAILURE PROBABILITY: COMPARISON OF BN AND MO METHODS

The posterior segment failure probability estimates from the BN and MO methods were compared. The method to estimate the posterior segment failure probability using the MO method was described briefly in Chapter 3, and in detail in Appendix F. Figs. 5.18 - 5.20 show the posterior failure probability estimates of the three segments, as a function of the number of cross sections representing the segment. The 95% confidence intervals around the BN estimate is also shown. Included in each figure is the MO estimate, which is given as a horizontal line, as the estimate is not a function of the number of cross sections in the BN. The agreement between methods is quite strong; in terms of reliability index, the differences are only 0.3%, 1%, and 1% for segments 13a, 13b, and 14, respectively.

Table 5.8 presents the prior and posterior segment failure probabilities estimated by the BN, for each of the segments, as well as for the system, for the case of a 1/400 year water level observation. The ratio of prior to posterior failure probability is presented, to show the reduction that results from making use of the survival observation. The ratio of prior to posterior system failure probability is 7.5, which means that the posterior failure probability is 7.5 times lower due to the survival observation. Table 5.9 shows the same results but for a 1/40 year water level observation. The impact on the failure probability is substantially less with the 1/40 year water level observation, with a ratio of prior to posterior system failure probability of only 2. An observation with a return period of 40 years is relatively high given the length of the record, but is not high enough to greatly impact a system with such a low prior failure probability.

POSTERIOR SYSTEM FAILURE PROBABILITY

The BN and MO/EP posterior system failure probability estimates were compared. The BN returned an estimate of 2.8E-4, and MO/EP method returned an estimate of 3.3E-4. This is a difference of about 18%. In terms of reliability indices, the estimates are 3.45 and 3.41 for the BN and MO/EP methods, respectively. This equates to a difference of about 1%.

Table 5.8: For a water level observation with a return period of 400 years (i.e. an extreme, rather unlikely observation): Prior and posterior segment failure probabilities for Segments 13a, 13b, and 14 computed with the BN and the MO method, and the system failure probability computed by the BN, and a combination of the MO and EP methods. The ratio of prior to posterior failure probability is also given.

$P(h > h_{obs}) = 1/400$						
Segment	BN prior	BN post	BN ratio	MO/EP prior	MO/EP post	MO/EP ratio
13a	6.8E-5	4.4E-5	1.6	9.0E-5	4.7E-5	1.9
13b	1.4E-3	1.2E-4	11.8	1.6E-3	1.4E-4	11.3
14	5.7E-4	1.5E-4	3.7	8.4E-4	1.8E-4	4.8
System	1.9E-3	2.8E-4	7.0	2.5E-3	3.3E-4	7.5

Table 5.9: For a water level observation with a return period of 40 years: Prior and posterior segment failure probabilities for Segments 13a, 13b, and 14 computed with the BN and the MO method, and the system failure probability computed by the BN and MO/EP methods. The ratio of prior to posterior failure probability is also given.

$P(h > h_{obs}) = 1/40$						
Segment	BN prior	BN post	BN ratio	MO/EP prior	MO/EP post	MO/EP ratio
13a	6.8E-5	6.6E-5	1.0	9.0E-5	6.9E-5	1.3
13b	1.4E-3	5.2E-4	2.7	1.6E-3	6.1E-4	2.6
14	5.7E-4	4.6E-4	1.2	8.4E-4	5.3E-4	1.6
System	1.9E-3	9.7E-4	2.0	2.5E-3	1.2E-3	2.1

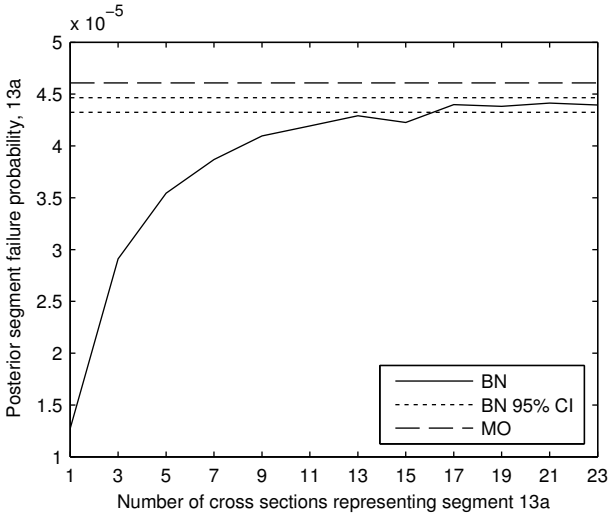

Figure 5.18: Posterior failure probability for segment 13a, computed with the BN and the MO method, using an observed water level with a return period of 400 years. The 95% confidence interval around the BN estimate is also presented.

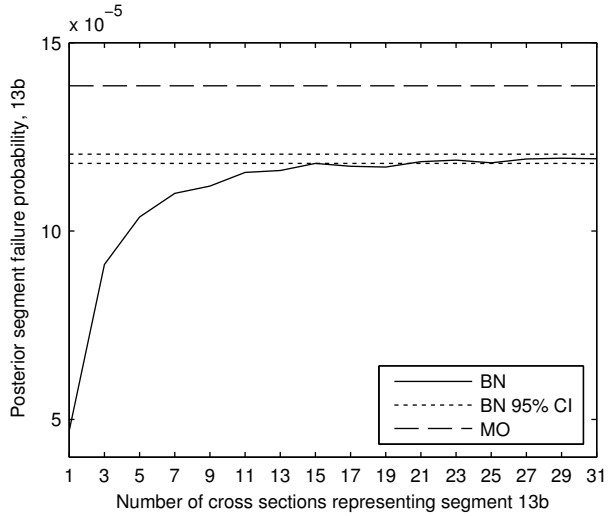

Figure 5.19: Posterior failure probability for segment 13b, computed with the BN and the MO method, using an observed water level with a return period of 400 years. The 95% confidence interval around the BN estimate is also presented.

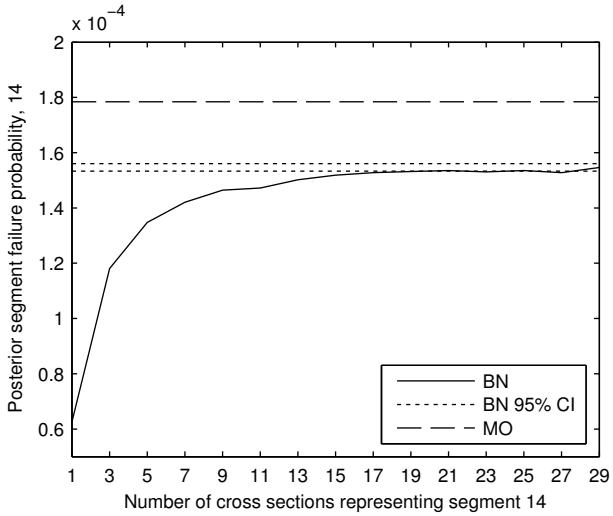

Figure 5.20: Posterior failure probability for segment 14, computed with the BN and the MO method, using an observed water level with a return period of 400 years. The 95% confidence interval around the BN estimate is also presented.

5.6. COMPUTATIONAL EFFICIENCY

In this case study, both the BN and the combined MO + EP methods from Hydra Ring were used to calculate the segment and system failure probabilities. This section focuses

on the time required for each of the methods to perform these calculations[1].

The BN has two calculation phases. The first phase is the iterative procedure to determine the necessary number of cross sections to represent the spatial variability of each levee segment. The second phase is the calculation of the system failure probability once the number of cross sections in each segment has been determined. The first phase is the most time consuming, and further research will look into ways to make it more efficient. In this chapter, the iterative procedure adds two cross sections at each iteration. This resulted in about 10 minutes to converge to the necessary number of cross sections, for each segment. Preliminary investigation into taking larger steps shows promise. The calculations in this chapter were repeated, but with three cross sections added per iteration (instead of two). The resulting segment failure probabilities were the same, but there was much faster convergence to the necessary number of cross sections, about 4 minutes per segment.

The second phase of the BN calculation, in which the number of segments in the system have already been determined, takes about 7 minutes. In the research in this dissertation, this was not prohibitively slow. Together with the first calculation phase, the entire analysis of prior and posterior segment and system failure probabilities takes about 20 minutes, if using an iterative procedure in the first phase in which three cross sections are added each iteration.

By comparison, the Hydra-Ring methods are much faster. Computing prior and posterior segment and system failure probabilities took about 4 seconds, more than half of which was used calculating the posterior influence coefficients (see Appendix F). This makes the Hydra-Ring algorithms about 300 times faster than the BN. Even excluding the computational costs of the iterative procedure (supposing this could be made substantially more efficient), the computational cost of the BN system calculation was still 7 minutes compared to 3 seconds, or about 100 times slower.

5.7. WHEN ARE SURVIVAL OBSERVATIONS USEFUL?

Survival observations are vastly available, but not always particularly useful. In this section I look at which observations are good candidates for a posterior analysis. Two primary factors determine how useful a survival observation will be: the magnitude of the prior conditional failure probability given the observed load, $P(\text{Fail}|S_{obs})$, and the relative influence that the resistance has on failure. Both of these factors will be described using examples.

The prior conditional failure probability, given the observed load, is a deciding factor in the usefulness of a survival observation. When the prior conditional failure probability is high, then observing survival is often very informative, because it tells us that our prior distributions are underestimating the strength of the levee. On the other hand, if the expected failure probability given the observed load is low, then observing survival essentially confirms our prior knowledge, and the posterior distributions are similar to the priors. This is illustrated with a simple example, where $Z = R - S$, and R and S are both normally distributed, $R \sim \mathcal{N}(\mu_R = 8, \sigma_R = 1)$ and $S \sim \mathcal{N}(\mu_S = 5, \sigma_S = 1)$. First, an observed load $S_{obs} = 7.5$ is considered, which corresponds to a return period of about 160

[1] Computation times are based on a 2.8 GHz computer with 8GB RAM

years, and results in a relatively high conditional failure probability $P(\text{Fail}|S_{obs}) = 0.3$. The top plot in Figure 5.21 shows the prior densities of R and S ($f_R(R)$ and $f_S(S)$), the observed load S_{obs} and the posterior or updated distribution of R (after incorporating the survival observation), $f_{Ru}(R)$. The bottom plot shows the prior and posterior conditional probability of failure, $P(\text{Fail}|S)$, as well as the density of S, $f_S(S)$. The total failure probability is given by Eq. 5.10, which shows that for a particular value of S, if either $f_S(S)$ or $P(\text{Fail}|S)$ are very low, the contribution of that load to the total failure probability will be low. Thus, the bulk of the failure probability comes from the range of S where both $f_S(S)$ and $P(\text{Fail}|S)$ are non-negligible (i.e. not too close to zero). When a survived load is observed, it affects the posterior curve $P(\text{Fail}|S)$; most notably, the curve will be zero for values of S less than S_{obs} [2] (see bottom plot in Figure 5.21). This means that the range of S below S_{obs}, where $f_S(S)$ is relatively high, will no longer contribute to the failure probability. This is the reason the failure probability decreases when a survived load is incorporated.

$$P(\text{Fail}) = \int \mathrm{P}(\text{Fail}|\mathrm{s})\,\mathrm{f_S}(\mathrm{s})\,\mathrm{ds} \tag{5.10}$$

Figure 5.22 shows the same example, but this time with an observed load of $S_{obs} = 6$, which corresponds to a return period of about 6 years, and results in a much lower conditional failure probability $P(\text{Fail}|S_{obs}) = 0.02$. The top plot in the figure shows that the updated density $f_{Ru}(R)$ is barely distinguishable from $f_R(R)$, and the bottom plot shows the posterior curve of $P(\text{Fail}|S)$ is barely distinguishable from the prior. In this case, removing the contribution of load values less than S_{obs} to the failure probability is fairly negligible.

Another deciding factor in the usefulness of a survival observation is the relative influence that the resistance has on failure. Intuitively this make sense, because if the resistance has little influence on the failure probability (i.e. the load is dominant), then improving its distribution will have little effect. Conversely, if the resistance is dominating the failure probability, then improving its distribution should have a large impact. Graphically, the influence of the resistance impacts the steepness of the conditional probability curve (less influence leads to a steeper curve, more influence leads to a broader curve). Via an example, I investigate how the relative influence of the resistance impacts the usefulness of a survival observation.

Consider a limit state function $Z = \beta + \alpha_R U_R + \alpha_S U_S$, where β is the prior reliability index, α_R and α_S are the influence coefficients of the (standard normally distributed) resistance variable U_R and load variable U_S. The influence coefficients are defined such that the square indicates the relative influence. For example $\alpha_R^2 = 0.3$ means that the resistance variable has a 30% influence on the failure probability. The sum of the squared influence coefficients is always 1, so that if $\alpha_R^2 = 0.3$, then $\alpha_S^2 = 0.7$. I considered several influence coefficients: $\alpha_R^2 = [0.9, 0.7, 0.5, 0.3, 0.1]$, three values of the reliability index: $\beta = [2, 3, 4]$, and two return periods of the observed load: 40 years and 400 years. The influence coefficient affects the shape of the conditional failure probability curve, and

[2] This assumes epistemic (fully reducible) uncertainty of the resistance variables. It is possible in some cases that there will still be some probability of failure for a value of the load below S_{obs}, but for simplicity this is not considered in the example.

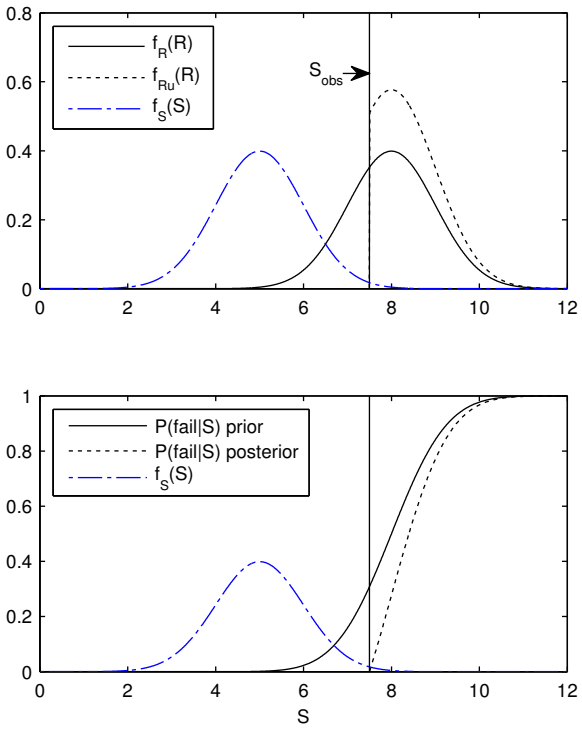

Figure 5.21: Top: Density of R and S, and the updated density of R, f_{Ru}, after incorporating the survival observation. Bottom: prior and posterior conditional failure probability curves, as a function of the load S, and the density of S, for an observed load with a return period of about **160 years**.

β and the observed load affect the area under the prior conditional failure probability curve that will no longer contribute to the posterior failure probability, after incorporating a survival observation.

Table 5.10 shows the reduction in the failure probability (as the ratio of prior to posterior) for an observed load with an extreme return period of 400 years. When the resistance is the dominant influence on the failure probability, the reduction in failure probability is substantial, even for a prior reliability index of $\beta = 4$. When the load is more dominant ($\alpha_S^2 = [0.7, 0.9]$, or equivalently $\alpha_R^2 = [0.3, 0.1]$), the reduction is much less, and for a $\beta = 4$ is completely negligible. Table 5.11 shows the reduction in failure probability for a load with a return period of 40 years, which is much more realistic to observe in practice than the 1/400 year load. In this case, when the resistance is the dominant influence on the failure probability, there is still a substantial reduction in the failure probability. With an influence of the resistance of 70% and a high prior reliability index $\beta = 4$, the posterior is 5 times lower than the prior. This is not an extreme reduction, but still quite useful for the small amount of effort required to perform the posterior analysis. However, once the influence of the resistance is 50% or lower, the reduction is negligible

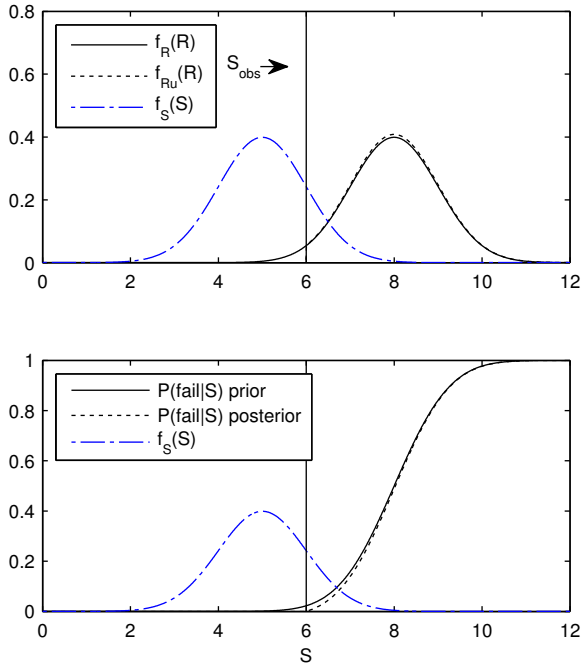

Figure 5.22: Top: Density of R and S, and the updated density of R, f_{Ru}, after incorporating the survival observation. Bottom: prior and posterior conditional failure probability curves, as a function of the load S, and the density of S, for an observed load with a return period of about **6 years**.

for $\beta = 4$, and quite small for $\beta = 3$.

Tables 5.10 and 5.11 show the reduction in failure probability for levees with different combinations of prior reliability, influence coefficients, and (extremity of) observed loads. Some of these combinations are seen more often than others, which is highlighted in this paragraph. I also provide the combination of these factors, and the reduction in failure probability for the primary levee case discussed in Section 5.5. In practice, particularly in the Netherlands, levees must be very reliable and comply with stringent limits on their failure probabilities. This means that in Tables 5.10 and 5.11, the $\beta = 4$ results are more representative of the prior failure probabilities encountered most often for primary levees. However, levees which fail to meet their safety standards are likely to have lower reliability indexes, more on the order of $\beta = 3$. Levees with a prior reliability index of $\beta = 2$ may be encountered for regional levees, or primary levees in countries with lower safety standards. Measurement records are usually between 50-100 years, which means observing the 1/400 year water level is not likely. Therefore, in practice, the results in Table 5.11 are more realistic than those in Table 5.10. The influence of the resistance will depend a great deal on the failure mechanism. For example, the resistance variables in macrostability have a large influence, whereas a mechanism like overtopping is dominated by the load. In Table 5.12, I present the combination of all of these factors (prior

Table 5.10: Reduction in failure probability (ratio of prior to posterior) due to incorporation of a survived load with a return period of **400 years**, for different influences of the resistance (10%, 30%, 70%, and 90%), and different prior reliability indexes β.

	$\alpha_R^2 = 0.9$	$\alpha_R^2 = 0.7$	$\alpha_R^2 = 0.5$	$\alpha_R^2 = 0.3$	$\alpha_R^2 = 0.1$
$\beta = 2$	2656	216	41	15	10
$\beta = 3$	1460	67	8	2	1
$\beta = 4$	839	23	2	1	1

Table 5.11: Reduction in failure probability (ratio of prior to posterior) due to incorporation of a survived load with a return period of **40 years**, for different influences of the resistance (10%, 30%, 70%, and 90%), and different prior reliability indexes β.

	$\alpha_R^2 = 0.9$	$\alpha_R^2 = 0.7$	$\alpha_R^2 = 0.5$	$\alpha_R^2 = 0.3$	$\alpha_R^2 = 0.1$
$\beta = 2$	245	26	6	2	1
$\beta = 3$	145	11	2	1	1
$\beta = 4$	91	5	1	1	1

reliability index, influence of the resistance variables, and extremity of the load) for the primary levee system piping example that was investigated in Section 5.5, including the reduction in failure probability (ratio of prior to posterior). The table shows that Segment 13b, which had the highest prior failure probability, also had the largest influence of the resistance variables, at about 60%. For both of these reasons, that segment had the largest reduction in failure probabability. For the other two segments, the influence of the resistance variables was about 40%. Notice that the reduction in failure probabilities in Table 5.12 (for the application to a real primary levee system) agree with the reductions provided for the generic example in Tables 5.10 and 5.11. One caveat to consider when referencing these tables is that if a real-life application includes resistance variables which are temporally variable, Tables 5.10 and 5.11 may overestimate the impact of a survival observation. This is because the impacts reported in them were calculated assuming that the uncertainty in resistance variables is fully reducible. However, in cases where the temporal variability of the resistance variables is considered negligible, Tables 5.10 and 5.11 can be used as guidance, to predict how useful a posterior analysis will be for a particular levee or levee system.

5.8. CONCLUSIONS

In this chapter, the BN methodology described in Chapter 2, including the incorporation of survival observations, was applied to two levee systems in The Netherlands. The first system was composed of regional levees protecting the city of Heerhugowaard from a canal. The prior failure probabilities due to piping were nearly 1, although no evidence of failure had been seen. Applying the BN to the three segments that made up this system, and imposing the survival observation (a high water level with no failure),

Table 5.12: Prior reliability index (β), influence coefficients of the load (α_S) and (collective) resistance (α_R) variables for the piping example presented in Section 5.5, and the reduction in failure probability (ratio of prior to posterior) for a 40-year water level observation (Red_{40}) and a 400-year observation (Red_{400})

Segment	Prior β	$\alpha_R^2 = 1 - \alpha_S^2$	Red_{40}	Red_{400}
13a	3.8	0.4	1.0	1.6
13b	3	0.6	2.7	11.8
14	3.3	0.4	1.2	3.7

reduced the failure probabilities by over two orders of magnitude. Discussions with the water board indicated that it was not necessarily poor prior distributions of the soil parameters, but rather an inappropriate choice of model (the Sellmeijer piping model) for the particular levee system. Specifically, they believe a mucky impermeable layer exists under the canal, which disconnects the pressure of the water level from the underlying sand layer. In this case, the prior distributions of the soil parameters may be correct, and the levee possibly quite weak; however, the disconnect makes it so that the load is not really felt, and therefore no failure occurs. Note that in this case, should the murky layer be disturbed, for example by dredging, the probability of failure could increase substantially.

The primary levee system that was considered, which protects the city of Zutphen from the IJssel River, was a more valid system for use of the Sellmeijer piping model, as it is unlikely that the IJssel river is underlain by an impermeable layer. The challenge with the primary system is that the prior probabilities are already quite low, so that a survival observation is not guaranteed to have much of an impact on the failure probability. For this case, I calculated posterior segment and system failure probabilities using two load observations (coupled with levee survival): a 1/40 year water level and a 1/400 year water level, to assess how the extremity of the observation influences the reduction in failure probability.

In both the regional and primary cases, I calculated prior and posterior cross-sectional, segment, and system failure probabilities. In the regional case, the system failure probability decreased by over two orders of magnitude, due to the extremely high prior failure probability of this system. In the primary case, the system failure probability decreased by 7.5 for a 1/400 year water level observation. Such an extreme water level is rarely observed, however. For a more realistic observation, the 1/40 year water level observation, the system failure probability only decreased by a factor of two. This confirms one of the disadvantages of using survival observations for primary defenses. The prior failure probability, conditional even on high loads, is already so low, that observing survival is not very informative.

For the primary case, I also calculated the prior and posterior segment failure probabilities using the MO method, which was described in Chapter 3, and in more detail in Appendices E and F. Chapter 3 had explored the validity of the MO method by investigating numerical examples. This chapter provides further verification in by comparing the results for this real-world application. Though the prior segment failure probabilities showed slightly worse agreement than the numerical examples in Chapter 3, they

were still in good agreement. In terms of reliability index, the MO prior estimates for segments 13a, 13b, and 14 were 2%, 1%, and 3% lower (note: lower reliability index = higher failure probability) than the BN estimates. The MO posterior estimates (for the case of a 1/400 year water level) were in even better agreement, a mere 0.3%, 1%, and 1% lower than the BN estimates. I further tested the agreement between the system failure probabilities estimated by the BN and the combined MO/EP methods, and found that in terms of reliability index, the differences were only about 1%. This difference is on par with the differences between the MO and BN segment estimates. I anticipated that the EP method would be fairly exact for a system of only three components, based on the results in Chapter 4. Still, it is a useful verification in an actual application, as opposed to the numerical examples explored in Chapter 4. The computation time for each method was looked at; the MO and EP methods required only about 4 seconds to compute the system failure probability, compared to about 20 minutes for the BN. The iterative procedure to find the number of cross sections to represent the segments in the BN is fairly time-consuming. Once the number of cross sections was determined, the system took about 7 minutes to compute. This is about 100 times slower than the MO and EP methods, though still not prohibitively slow.

This chapter included a section investigating under which conditions survival observations are useful. In general, they are most useful (i.e. they result in a large reduction in the failure probability estimate) when the prior failure probability of the levee (system) conditional on the observed load is relatively high, and/or when the relative influence of the resistance variables on the failure probability is high. The calculations presented in this chapter show that for a 1/40-year water level observation, a survival observation will result in a large reduction of the failure probability if the prior failure probability is relatively high ($\beta = 2$), even when the resistance is not too influential. For lower prior failure probabilities, reductions can be expected only if the resistance is fairly influential. For example, for a prior reliability of $\beta = 4$, only a relative influence of 70% or more results in a reduction of the failure probability. Tables 5.11 and 5.10, which show the reduction in failure probability due to a survival observation for different values of the prior failure probability and influence coefficients, are a good guidance for anyone considering updating levee failure probability with survival observations. Consulting these tables prior to carrying out reliability updating can help avoid unnecessary calculations.

6

ESTIMATING GEOTECHNICAL FAILURE MODEL UNCERTAINTY

6.1. INTRODUCTION

Geotechnical failure models are essential tools in levee reliability modeling. They estimate the resistance of a levee to a given load, with respect to a specific type of failure, such as piping, slope instability, or erosion. Failure models are prone to error because they try to capture complex geotechnical processes using simplified approximations, especially the computationally-efficient models that are well-suited to reliability analysis. Accurately quantifying the error (or uncertainty) in the model output is important because it can have a substantial influence on the estimated failure probability of the levee.

Estimating the uncertainty in failure model output is challenging. Validating such a model with field experiments is costly and not always feasible. And even with a field test, it is still uncertain if the model would perform the same at a different site or under different conditions, making it necessary to carry out multiple field experiments at different locations. Expert opinion is another method of quantifying the uncertainty in failure model output. However, because experts have seldom observed failures, it is questionable how well they can estimate this uncertainty. Because the model uncertainty can have a large influence on the estimated failure probability of a levee, it is preferable to estimate the model uncertainty distribution in a reproducible, quantitative way.

In this chapter, I propose a method to estimate the uncertainty in a failure model with a Bayesian network (BN) using failure/survival observations at locations where the load event was recorded, and where there is sufficient site/subsurface data (including, but not limited to field experiments). The BN developed in this chapter is used to estimate the posterior (updated) distributions of the model uncertainty parameters. In the research presented here, a lognormal distribution of the model uncertainty is assumed, which has two parameters (μ and σ), which means our posterior is only two-dimensional. I exploit the low dimensionality of the problem by using numerical integration to perform inference. This is advantageous because it is one of the most accurate probabilistic methods,

and in the case of two dimensions, also incredibly fast.

This chapter is laid out as follows. Section 6.2 presents the methodology to calculate the model uncertainty based on observations and data, using a BN. Section 6.3 uses a synthetic example to explore the impact of different choices/constraints in the methodology, such as the importance of the prior distributions, and the uncertainty due to limited data. Section 6.4 applies the method to a case study, focusing on the slope stability model D-Geostability, applied to a number of locations in the Netherlands. The resulting posterior model uncertainty distribution from this application was used to support the development of calculation rules for the Dutch levee assessments.

6.2. METHODOLOGY

This section describes the development of a BN for estimating failure model uncertainty based on hindcasts and performance observations. A hindcast is the output of a failure model for (sharply estimated) input from past events. The idea is that given the past conditions as input, the model should predict the performance observation (failure or survival). The method described herein is applicable whenever performance of the levee has been observed for a known loading event at multiple locations. The analysis should only include locations where the levee geometry and the subsurface have been explored in enough detail that the model input is considered certain. This is important so that the calculated uncertainty in the model output will be due to the error in the model, and not to error/uncertainty in the input. In practice, there will always be some uncertainty in the input, but the goal is to keep it low relative to the uncertainty in the model. Section 6.2.1 discusses how to characterize failure model uncertainty, Section 6.2.2 describes how to represent the problem of an unknown failure model uncertainty distribution in a BN, and how to interpret failure observations. This latter has to do with the discrepancy between a 'failure' observation and the output of a failure model (which is usually a comparison of the resistance to the load). The way that failure is interpreted is relevant because it has consequences for the posterior probability of the model uncertainty parameters. Section 6.2.3 describes how inference is performed using survival/failure observations and failure model output, and Section 6.2.4 discusses the choice of prior distributions.

6.2.1. CHARACTERIZATION OF FAILURE MODEL UNCERTAINTY

The error in a failure model can be complex, but it is commonly characterized as either an additive error or a multiplicative one. Consider a failure model that predicts the stability factor $G_m(\mathbf{X})$, which is the ratio of the modeled resistance $R_m(\mathbf{X})$ to the modeled load $S_m(\mathbf{X})$. The input to the failure model is a set of (potentially random) variables $\mathbf{X} = [X_1, X_2, ..., X_n]$ (e.g. soil porosity, aquifer thickness, etc.). The subscript m refers to the fact that these quantities are modeled. In the remainder of the chapter $G_m(\mathbf{X})$ is shortened to G_m to make the equations and text cleaner, but the dependence on the input variables remains.

$$G_m(\mathbf{X}) = \frac{R_m(\mathbf{X})}{S_m(\mathbf{X})} \tag{6.1}$$

The stability factor is describing the brink between failure and survival of a structure.

The 'true' stability factor G is the ratio of the true resistance to the true load $G = R/S$ (the absence of the subscript m indicates true quantities). $G = 1$ represents the limit state between survival and failure; when $G > 1$ the structure survives and when $G < 1$, the structure will fail. This is not the case with G_m because the model may be biased, and will be uncertain due to the inexact modeling of the physical processes involved in geotechnical failure. Consider failure to be conditional upon the value of the stability factor. For the true stability factor, this conditional failure probability is a step function (see Equation 6.2). For G_m this will not be the case, because the error in the model makes it possible that $G_m < 1$ in cases where the structure survives, or conversely $G_m > 1$ when the structure fails.

$$P\left(\text{Fail}|G = g\right) = \begin{cases} 0, & g > 1 \\ 1, & g \leq 1 \end{cases} \tag{6.2}$$

The uncertainty in the model (denoted as M) can be described as a multiplicative error or an additive one. Ideally, measurements would determine whether a multiplicative or additive representation is better suited. Errors that depend on the magnitude of the predicted stability factor would be better represented by a multiplicative error, and errors that are independent of the prediction would be better represented by additive error. In the case of predicting levee failures, measured errors are hard to come by, and there are definitely too few to determine this conclusively. In practice, a multiplicative lognormal distribution is a common choice for the model uncertainty ([76]). The advantage of such a representation is that the stability factor is guaranteed to always be positive, which is necessary as negative stability factors have no physical meaning. There are ways to accomplish this with an additive error, for example by truncating the uncertainty distribution, but it is less elegant and has no obvious advantage. Because the range of calculated stability factors is relatively constrained around 1, choosing an additive or multiplicative error is unlikely to make a substantial difference. In this chapter I focus on a multiplicative lognormal distribution for the model uncertainty. In this case, the true (unknown) stability factor is represented as the product of the model uncertainty and the modeled stability factor; see Eq. 6.3. This distribution of M has two parameters, μ and σ, which are the mean and standard deviation of $\ln(M)$. The density of M is given in Eq. 6.4.

$$G = M \cdot G_m \tag{6.3}$$

$$f_M(m) = \frac{1}{m\sigma\sqrt{2\pi}} \exp\left[-\frac{\left(\ln m - \mu\right)^2}{\sqrt{2}\sigma}\right] \tag{6.4}$$

The value of the parameters μ and σ are unknown. In the following section I present a BN which uses observations and hindcasts to estimate the posterior distribution of these parameters, and (by extension) the posterior distribution of M.

6.2.2. BAYESIAN NETWORK FOR MODEL UNCERTAINTY ESTIMATION

The BN to estimate the model uncertainty parameters μ and σ is illustrated in Figure 6.1. The nodes in the BN are the parameters of the model uncertainty distribution, μ and σ,

the model uncertainty M_i, the (true) stability factors G_i, the modeled stability factors G_{mi}, and the observed failure states of the levee $Fail_i$, for the $i = 1...n$ locations. The nodes M_i are lognormally distributed with parameters μ and σ. The stability factors G_i are a function of M_i and $G_{m,i}$: $G_i = M_i \cdot G_{m,i}$. The node $Fail_i$ is also a function, it is equal to 1 when $G_i < 1$ and 0 when $G_i > 1$. Functional nodes are indicated with black edges. The modeled stability factors are deterministic, which are represented by square nodes[1].

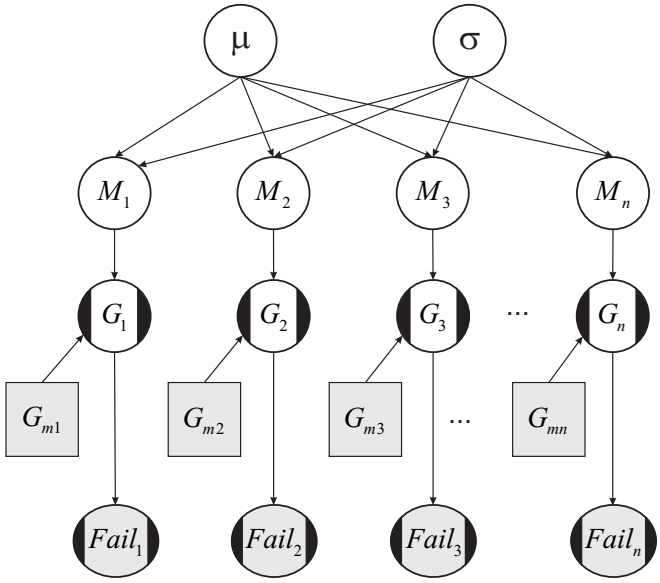

Figure 6.1: BN for estimating the lognormal model-uncertainty parameters μ and σ based on modeled stability factors G_{mi} and observed levee performance ($Fail_i = 0$ or $Fail_i = 1$). Circular nodes represent random variables, black edges represent functional nodes, square nodes represent deterministic values, and grayed nodes indicate observed variables.

Observed nodes, or nodes with known values, are indicated by a gray fill in Figure 6.1. Often $Fail_i$ is observed, but not the exact value of the stability factor G_i. This is because most data is historical, and the exact load at which failure occurred is generally unknown. Still, observing $Fail_i$ and knowing $G_{m,i}$ gives us information about G_i and M_i (via the relationship Eq. 6.3 and Eq. 6.5). These two pieces of information tell us that when $Fail_i = 1$, $G_i = M_i \cdot G_{m,i} < 1$, and therefore that $M_i < 1/G_m$.

$$\begin{aligned} G_i < 1, \quad & Fail_i = 1 \\ G_i > 1, \quad & Fail_i = 0 \end{aligned} \tag{6.5}$$

During field tests, it is possible to observe $G_i = 1$, which occurs at the load which *just* exceeds the resistance. In such a case, the observation can be considered to be $M_i = 1/G_{m,i}$. It is also possible with field tests to consider the observation of G to be a range, due to other potential sources of uncertainty in the model input, or related to the test site. In that case, G can be considered to be observed between some lower bound

[1] In BNs, square nodes can refer to decision nodes; in this chapter they only refer to deterministic nodes.

G_{LB} and 1: $G_{LB} < G < 1$, which translates to $G_{LB}/G_m < M < 1/G_m$. In the case of survival, when $Fail_i = 0$, it is only known that $G_i > 1$, and that $M_i > 1/G_m$. The relationship between observed variables and M are summarized in Table 6.1.

Table 6.1: Information about the model uncertainty M derived from observations of *Fail* and (if observed) G

Fail	G	M
1	*not observed*	$M < 1/G_m$
1	$G = 1$	$M = 1/G_m$
1	$G_{LB} < G < 1$	$G_{LB}/G_m < M < 1/G_m$
0	*not observed*	$M > 1/G_m$

The joint distribution over all the nodes in the network is given in Eq. 6.6.

$$P(\mu,\sigma,M_1,\ldots,M_n,G_1,\ldots,G_n,\text{Fail}_1,\ldots\text{Fail}_n) = P(\mu)\,P(\sigma)\prod_{i=1}^{n} P(M_i|\mu,\sigma)P(G_i|M_i)\,P(Fail_i|G_i) \tag{6.6}$$

The objective is to estimate the posterior distribution of μ and σ, given the observed variables in the network. From Bayes theorem, the posterior distribution can be written according to Eq. 6.7.

$$P(\mu,\sigma|\mathbf{G_m},\mathbf{Fail}) \propto P(\mu)\,P(\sigma)\,P(\mathbf{G_m},\mathbf{Fail}|\mu,\sigma) \tag{6.7}$$

The third term is the likelihood of the observations, given the parameters μ and σ. It is the product over the likelihoods at the individual locations; see Eq. 6.8. Because the observations of $G_{m,i}$ and $Fail_i$ translate into information about M_i, the individual likelihoods can be written in terms of M. The form of the likelihood will depend on which data were observed. When only $G_{m,i}$ and $Fail_i$ are observed (as illustrated in Fig. 6.1), then failure only tells us that $G_i < 1$. In this case, the likelihood is given in Eq. 6.9, where F_m is the cumulative distribution function of M, which is taken to be the lognormal distribution in this chapter.

$$P(\mathbf{G_m},\mathbf{Fail}|\mu,\sigma) = \prod_{i=1}^{n} P(G_{m,i},\text{Fail}_i|\mu,\sigma) \tag{6.8}$$

$$P(G_{m,i},\text{Fail}_i|\mu,\sigma) = \begin{cases} P(M_i < 1/G_m|\mu,\sigma) = F_M(1/G_m;\mu,\sigma), & Fail_i = 1 \\ P(M_i > 1/G_m|\mu,\sigma) = 1 - F_M(1/G_m;\mu,\sigma), & Fail_i = 0 \end{cases} \tag{6.9}$$

If $G_i = 1$ is observed in the case of failure, this is a stronger observation than just observing $Fail_i$. The likelihood of the observation for this case is given in Eq. 6.10. The observation of G_i translates into $M_i = 1/G_{m,i}$, and the probability density f_M is used in the case of failure, instead of the distribution function F_m.

$$P(G_{m,i},\text{Fail}_i|\mu,\sigma) = \begin{cases} P(M_i = 1/G_m|\mu,\sigma) = f_M(1/G_m;\mu,\sigma), & Fail_i = 1 \\ P(M_i > 1/G_m|\mu,\sigma) = 1 - F_M(1/G_m;\mu,\sigma), & Fail_i = 0 \end{cases} \tag{6.10}$$

Sometimes $G_i = 1$ can be an over-confident estimate for failure during a field test, for example due to uncertainties in the input or issues related to how well the test site mimics reality. Still, it is likely to be close to 1. In these cases, experts can impose a lower bound $G_{LB,i}$ on what G_i can be, given failure. That means $G_{LB,i} < G_i < 1$, or $G_{LB,i}/G_{m,i} < M_i < 1/G_{m,i}$ (see Table 6.1). The likelihood for this case is given in Eq. 6.11. It takes the likelihood to be the area under the density of M between G_{LB} and 1.

$$P\left(G_{m,i}, \text{Fail}_i | \mu, \sigma\right) = \begin{cases} F_M\left(1/G_{m,i}; \mu, \sigma\right) - F_M\left(G_{LB,i}/G_{m,i}; \mu, \sigma\right), & Fail_i = 1 \\ 1 - F_M\left(1/G_{m,i}; \mu, \sigma\right), & Fail_i = 0 \end{cases} \tag{6.11}$$

In the following section I will describe methods to sample and perform inference in the BN, to derive the posterior distributions of μ and σ, and ultimately, of M.

6.2.3. SAMPLING AND INFERENCE IN THE BAYESIAN NETWORK

Many inference methods are possible to solve Eq. 6.7, including MC rejection sampling, MC importance resampling ([77]), and Markov Chain MC (MCMC). However, because the posterior in this case is two-dimensional, I opted for numerical integration as it is highly accurate, and very fast in two dimensions.

Numerical integration considers a two-dimensional grid over possible μ-σ values. For each grid cell, the posterior distribution is calculated according to Eq. 6.12. The denominator in Eq. 6.12 is a normalizing constant that ensures the total probability is equal to 1 (hence the equal sign in Eq. 6.12 instead of the $\propto$ in Eq. 6.7). The densities f_μ and f_σ are the prior densities of μ and σ.

$$P\left(\mu_j, \sigma_k | \mathbf{G}_m, \mathbf{Fail}\right) = \frac{f_\mu\left(\mu_j\right) \cdot f_\sigma\left(\sigma_k\right) \cdot \prod_{i=1}^{n} P\left(G_{m,i}, \text{Fail}_i | \mu_j, \sigma_k\right)}{\sum_j \sum_k f_\mu\left(\mu_j\right) \cdot f_\sigma\left(\sigma_k\right) \cdot \prod_{i=1}^{n} P\left(G_{m,i}, \text{Fail}_i | \mu_j, \sigma_k\right)} \tag{6.12}$$

The posterior distributions of μ and σ in Eq. 6.12 are used to obtain the posterior distribution of M. To do this, a range of M is discretized, say $[M_{LB}, M_{UB}]$, where M_{LB} and M_{UB} are values low and high enough (respectively) that their cumulative probability will be 0 and 1 (respectively). For each value of m_l within the range of M, the associated posterior non-exceedance probability is computed according to Eq. 6.13.

$$P\left(M < m_l\right) = \sum_j \sum_k F_m\left(m_l; \mu_j, \sigma_k\right) P\left(\mu_j, \sigma_k | \mathbf{G}_m, \mathbf{Fail}\right) \tag{6.13}$$

The parameters that can be adjusted in using NI are the fineness of the μ-σ grid, as well as the fineness of the discretization of M, when computing its posterior probability.

6.2.4. CHOICE OF PRIOR DISTRIBUTION

The prior distributions of the model uncertainty parameters will influence their posterior distributions (see Eq. 6.12). As the amount of data increases (i.e. number of locations with hindcasts and performance observations), the influence of the priors will decrease. When little is known about the distribution parameters, wide priors - such as uniform

distributions - can be used. Generally there is some intuition about the parameters. If not the values of the parameters themselves, then at least intuition about the mean and variance of the distribution of M, which can be used to derive more narrow, informed priors of μ and σ. The mean and standard deviation (S.D.) of M are related to μ and σ according to Eqs. 6.14 and 6.15.

$$\text{Mean} = \exp\left(\mu + \frac{\sigma^2}{2}\right) \tag{6.14}$$

$$\text{S.D.} = \sqrt{\left(\exp\left(\sigma^2\right) - 1\right)\exp\left(2\mu + \sigma^2\right)} \tag{6.15}$$

Expert intuition about realistic lower and upper bounds on M can be used to constrain the priors. This is discussed in more detail in the synthetic example, in Section 6.3.2. A cross-validation can help determine which prior results in the most predictive posterior. Details about cross-validation are provided in the application section, see Section 6.3.2.

6.3. SYNTHETIC EXAMPLE

This section illustrates and explores the method via a synthetic example, in which μ and σ, the distribution parameters of the lognormally distributed model uncertainty M, are known. I investigate the influence that the choice of prior distribution for μ and σ and the number of data points have on the posterior predictive distribution of M. In the first part of the example, sections 6.3.1 - 6.3.4, I only consider the case where *Fail* is interpreted as $G < 1$. This is representative of historical data, where the exact load at the time of failure is unknown. In section 6.3.6, I modify the example and consider the interpretation of failure $G_{LB} < G < 1$, in addition to $G < 1$, and assess the influence it has on the posterior predictive distribution of M. The synthetic data in this example are modified to be more representative of field experiment cases. Specifically the generated values of G in Section 6.3.6 are closer to 1 than in the synthetic example considered in sections 6.3.1 - 6.3.4.

6.3.1. SYNTHETIC DATA

Synthetic data is generated by defining a known 'true' distribution of the model uncertainty M. True values are assigned to the parameters μ and σ, $\mu_T = 0.6$ and $\sigma_T = 0.2$, where the subscript T indicates 'true'. The parameters of the lognormal distribution are the mean and standard deviation of $\ln(M)$, not M. The associated 'true' mean and standard deviation of M are 1.9 and 0.4, respectively. The subscript T is retained in this synthetic example to distinguish between the known (assigned) true values of the parameters of M (μ_T and σ_T) and the prior and posterior estimates of μ and σ. The model uncertainty distribution is presented in Figure 6.2.

A set of 'true' stability factors G_T are generated by sampling from a uniform distribution with specified upper and lower bounds: $G_T \sim \mathcal{U}(a,b)$, where $a = 0.3$ and $b = 1.7$. The choices for the bounds are somewhat arbitrary, but these represent a realistic range. Again, the subscript T for true is retained to distinguish the generated values of the stability factors G_T from the estimated distribution of the stability factor $G = M \cdot G_m$.

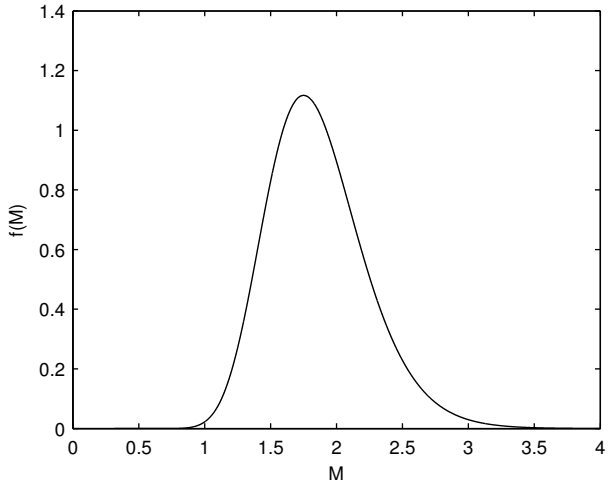

Figure 6.2: Assigned 'true' distribution of the model uncertainty M for the synthetic example.

The synthetic modeled stability factors $G_{m,i}$ are then calculated as $G_{m,i} = G_T/m_i$, where m_i is a sample from the distribution M. The synthetic failure observations are directly derived from the values of G_T: $Fail = 0$ when $G_T > 1$ and $Fail = 1$ when $G_T < 1$.

Synthetic datasets of differing lengths (n) were considered, n =12, 25, 50, and 100. For each length, I considered the relevance of the prior on the goodness of the posterior estimates (in general, the fewer data, the more influential the priors). I also investigated how the length of the dataset contributed to the variance in the quantiles of the posterior predictive distribution of M.

6.3.2. PRIORS

Four priors were considered: an informed and uninformed, both with and without expert constraints. The informed prior is one in which experts have some belief about the distributions of μ and σ, and the uninformed prior is one in which experts can only guess at the minimum and maximum values for the two parameters. The expert constraints reflect intuition that experts have about minimum and maximum values of specific quantiles of M. For example, an expert may believe that $P(M < m^*) = 0.01$. This means the expert believes there is only a 1% probability that M can be lower than the quantile m^*.

For the informed prior a normal distribution for μ was assumed (which can take both positive and negative values), and a lognormal distribution for σ (which can only take positive values). For the uninformed prior, it was assumed that the experts have no intuition about the parameters μ and σ except for bounds on their realistic values. In the synthetic example, the uninformed prior was taken to be uniform for both parameters.

The expert constraints provide information about M, which in turn constrains the values of μ and σ. The expert constraints for this synthetic example were assigned to be $P(M < 0.5) = 0.01$ and $P(M > 4) = 0.01$. The way of implementing the constrained priors was to first independently sample the priors of μ and σ, and then essentially discard μ-σ

samples for which $M(\mu,\sigma)$ does not comply with the constraints, by assigning a likelihood of zero to these samples. The four prior cases considered are summarized in the bulleted list below.

- Case 1: Uniform priors:
  $\mu \sim \mathcal{U}(a_\mu, b_\mu)$ and $\sigma \sim \mathcal{U}(a_\sigma, b_\sigma)$

a_μ	b_μ	a_σ	b_σ
-1.5	1.5	0	0.5

- Case 2: Informed priors:
  $\mu \sim \mathcal{N}\left(\mu_\mu, \sigma_\mu^2\right)$, $\sigma \sim \ln\mathcal{N}\left(\mu_\sigma, \sigma_\sigma^2\right)$

μ_μ	σ_μ	μ_σ	σ_σ
0	0.5	0.1	0.1

- Case 3: Uniform priors with expert constraint:
  Same as Case 1, but imposing a likelihood of zero for any μ-σ sample in which $P(M < M_{LB}) > P_{LB}$ or $P(M > M_{UB}) > P_{UB}$.

M_{LB}	P_{LB}	M_{UB}	P_{UB}
0.5	0.01	4.0	0.01

- Case 4: Informed priors with expert constraint:
  Same as Case 2, but imposing a likelihood of zero for any μ-σ sample in which $P(M < M_{LB}) > P_{LB}$ or $P(M > M_{UB}) > P_{UB}$.

Figures 6.3 and 6.4 show the four priors for the parameters μ and σ, respectively.

6.3.3. SIMULATIONS

Posterior predictive distributions of the model uncertainty M were calculated using synthetic data generated as described in Section 6.3.1, for the following combinations:

- Four priors: Uniform, Informed, Uniform with expert constraints, Informed with expert constraints
- Four lengths of data: $n = 12, 25, 50, 100$

In total, this leads to $4 \times 4 = 16$ combinations. In addition, each combination was calculated 100 times, to estimate the variance in the posterior predictive distribution of M that is due to the length of the data set.

A 100×100 grid was chosen for the numerical integration, with μ values ranging from -1.5 to 1.5, and σ values ranging from 0.01 to 0.5. For the derivation of the posterior predictive M distribution (see Eq. 6.13), M was discretized into 1000 values between 0 and 10.

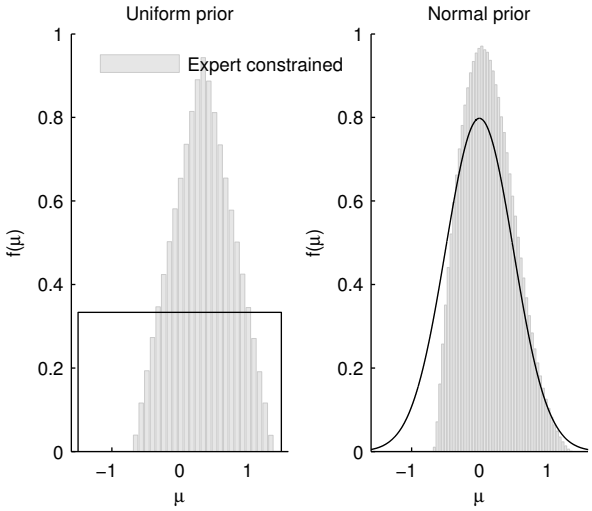

Figure 6.3: Prior distributions of μ

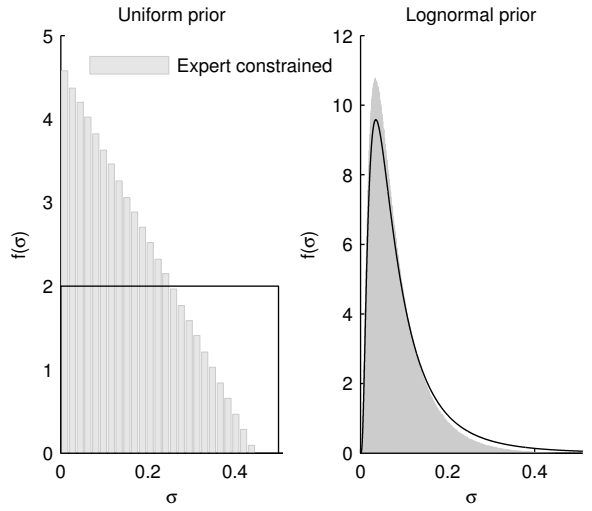

Figure 6.4: Prior distributions of σ

6.3.4. RESULTS

Each combination of prior distributions of μ and σ and length of the data set, n, led to a different posterior predictive distribution for M. An insightful way to present the results is to compare the quantiles of the posterior predictive distribution of M with the quantiles of the known (true) distribution, for each of the cases considered.

QUANTILES OF M

The quantiles of the posterior predictive M distribution are presented using box and whisker plots. For each case of prior distributions of μ and σ and length of the data set, n, the analysis was run 100 times, generating different data sets each time, to estimate the variance in the posterior predictive M distribution that results from the limited amount of data. In the plots, the bottom and top of the box represent the 25% and 75% quantile (denoted q_1 and q_3), respectively. The whiskers extend $1.5 \times (q_3 - q_1)$ from the top and bottom of the box. For reference, this is +/- 2.7 standard deviations for normally distributed data. The horizontal line in the middle of the box shows the median. When that line is not in the center, it indicates a skewed distribution. Any calculated quantiles of M which fall outside the whiskers are shown with an 'x' symbol, and represent outliers. Each figure contains four subplots, one for each of the data lengths that were considered.

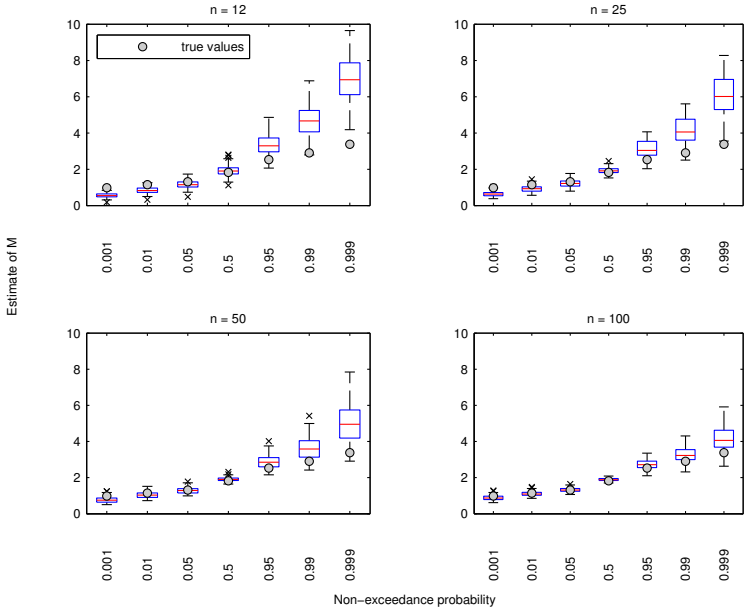

Figure 6.5: Quantiles of the posterior predictive distribution M, for the case of uniform priors, for n =12, 25, 50, and 100 observations.

For all priors, as the number of data points increases, the variance in the quantiles of M decrease. The greatest variance occurs in the upper tails of the posterior predictive M distribution (the 99.9% quantile). However, the lower tail is usually of more interest.

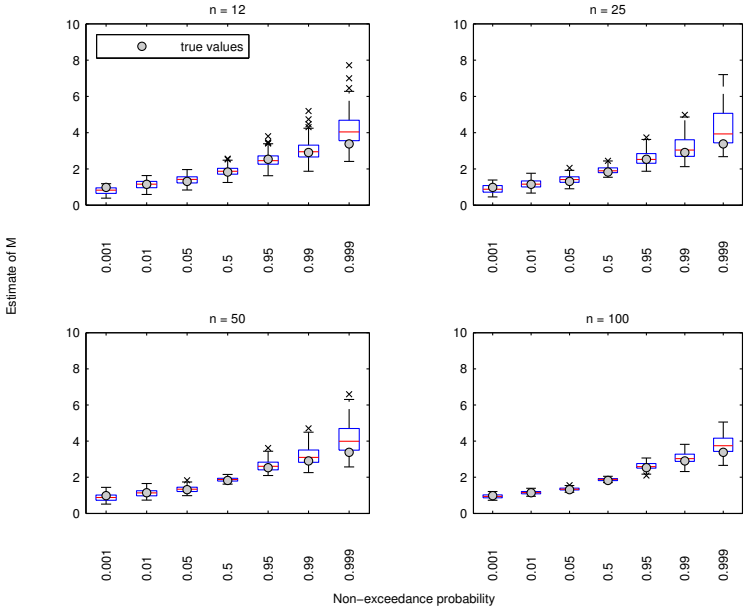

Figure 6.6: Quantiles of the posterior predictive distribution M, for the case of informed priors, for n =12, 25, 50, and 100 observations.

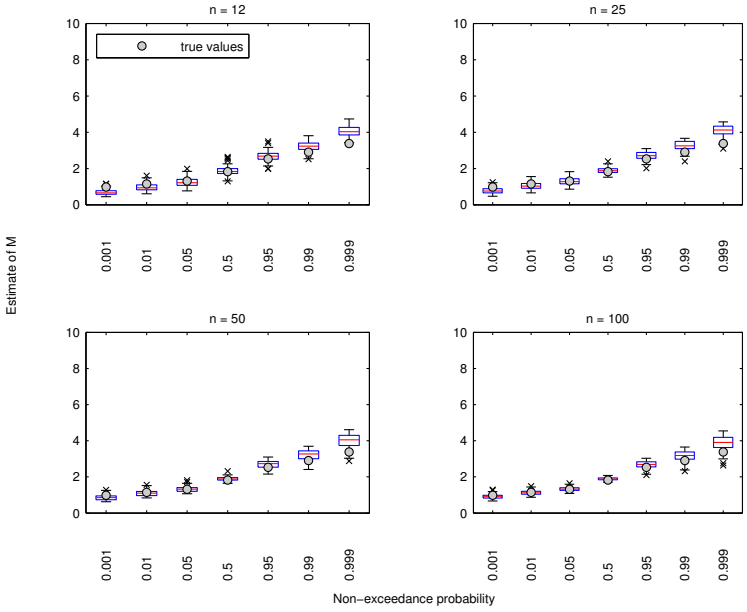

Figure 6.7: Quantiles of the posterior predictive distribution M, for the case of uniform priors with expert constraint, for n =12, 25, 50, and 100 observations.

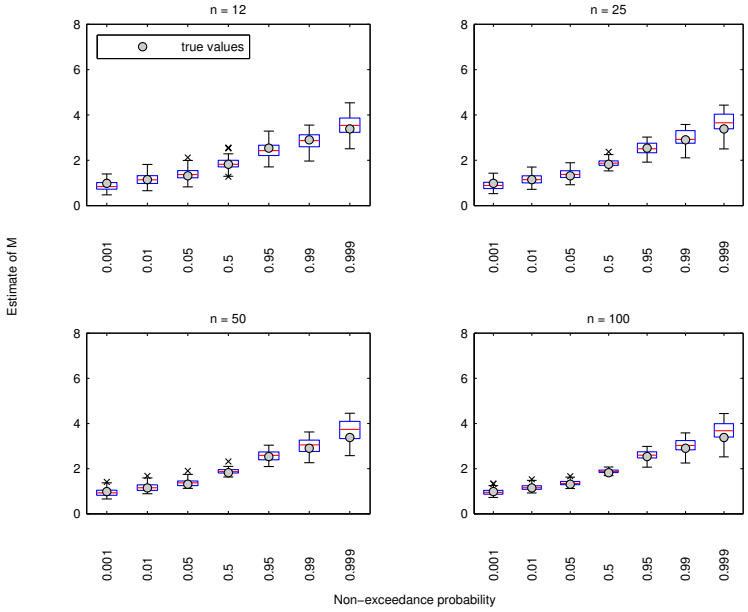

Figure 6.8: Quantiles of the posterior predictive distribution M, for the case of informed priors with expert constraint, for n =12, 25, 50, and 100 observations.

For example, in Dutch flood defense assessments, the 5% quantile is used in a semi-probabilistic approach (where the limit state is computed deterministically using characteristic values), as a conservative model factor. Because NI is a fairly exact method, the variance can be considered to be entirely due to the limited data.

The 'informed' priors - a normally distributed μ and a lognormally distributed σ - result in less variance, and better estimates of the quantiles, than the uniform priors. This is because higher values of the variance in M are excluded by the informed priors, which restricts those unrealistically high quantiles in the tail of the posterior predictive M distribution (see Figure 6.4).

In the case of expert-constrained priors, the results are much better. This is because incorporating knowledge about realistic upper and lower bounds on M strongly reduces the variance in high and low quantiles of M. Both the uniform and informed priors estimated the posterior predictive distribution of M very well when coupled with expert constraints, even with limited data sets.

6.3.5. AVOIDING BIAS: CHOICE OF HISTORIC OBSERVATIONS

In the previous synthetic example, the 'true' stability factors G_T were uniformly generated within the range 0.3 to 1.7 (see Section 6.3.1) in such a way that approximately half were failures and half were survivals. In reality there tend to be more data about the site conditions when a failure occurred. So even though survival occurs far more often than failure, there are more failure observations that are suited for this type of analysis than

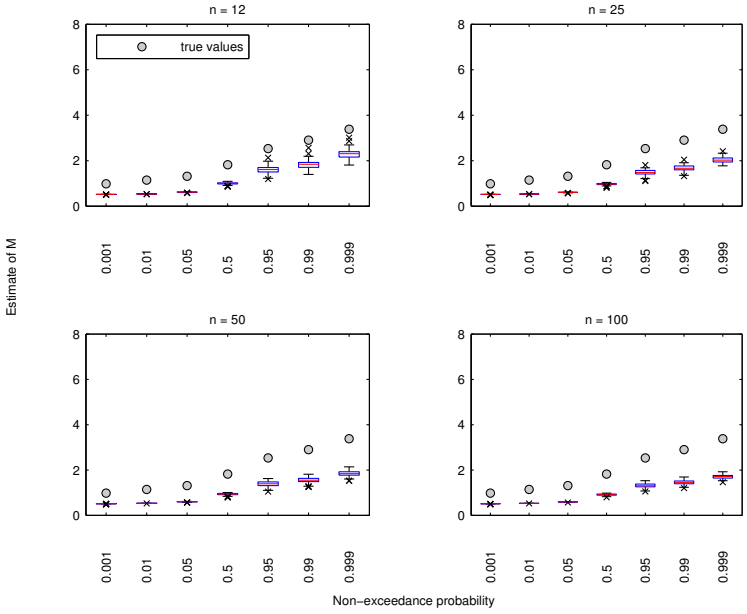

Figure 6.9: Bias introduced by only considering failure observations. Quantiles of the model uncertainty M estimated by NI, for the case of informed priors with expert constraint, for n =12, 25, 50, and 100 observations.

survival observations. Selecting observations of all one type - either failure or survival - will lead to bias in the estimated quantiles of the model error distribution. To illustrate this, the synthetic example was rerun for the scenario in which only failures were observed. In the interest of brevity, only one prior case was considered (informed priors constrained by expert opinion, Case 4 in Section 6.3.2). The exclusion of survival observations led to the quantile estimates shown in Figure 6.9. Compared with the results in Figure 6.8 (in which roughly equal numbers of failures and survivals were considered), a bias is clearly observed.

The synthetic example was further used to explore what percentage of the observations must be survivals to avoid a bias in the quantile estimates of the model error distribution. If at least 20% survival observations are included, the bias seems mostly eradicated, as shown in Figure 6.10. It would not be prudent to assume this percentage will be valid in all real life applications, but it does give the impression that equal numbers of survival and failure observations are not required. Quality is probably also more important than quantity. Specifically, observations that go against expectations (i.e. model predictions) will likely be the most informative for estimating the model uncertainty distribution. Note further that most historic survivals relate to relatively low or moderate loading levels which have little influence on the result. Important realizations of the model uncertainty may be missed if survivals during high load conditions are not included in the data set. This should be kept in mind when constructing the data set.

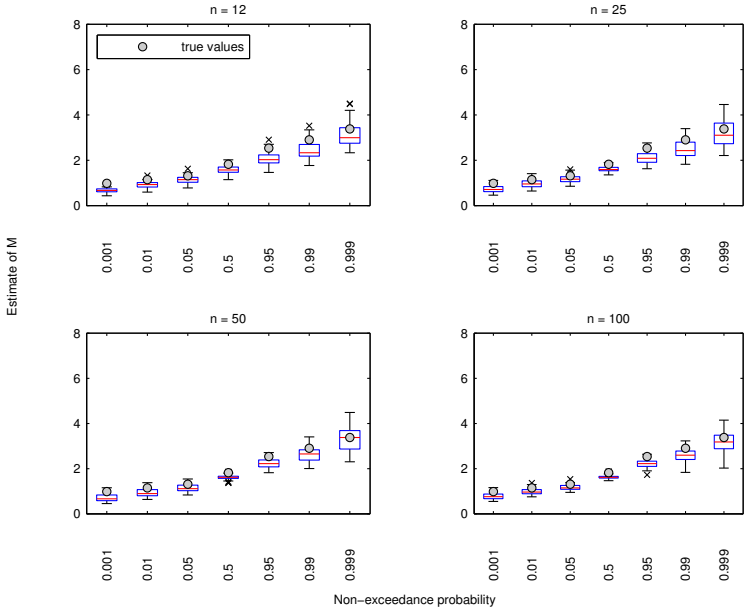

Figure 6.10: Considering 80% failure observations. Quantiles of the model uncertainty M estimated by NI, for the case of informed priors with expert constraint, for n =12, 25, 50, and 100 observations.

6.3.6. INTERPRETING FAILURE OBSERVATIONS

In this section, the interpretation that failure indicates $G_{LB} < G < 1$ is compared with the interpretation that failure indicates $G < 1$, specifically for field experiments, where the load at which failure occurred is known. In Section 6.3.1, it was described that the 'true' stability factors G_T were generated, $G_T \sim \mathcal{U}(a,b)$, where $a = 0.3$ and $b = 1.7$. In the current section, stability factors for failure are generated that are closer to 1, to better represent the situation that failure - and the load at which it occurred - were observed in the field. The density of the (generated) 'true' stability factors used in this section is given in Eq. 6.16, where $G_{max} = 1.7$ as in the previous example. The choice of G_{LB} will be described below.

$$f_{G_T}(G_T) = \begin{cases} 0.5 \cdot \frac{1}{(1-G_{LB})} & G_T \in [G_{LB}, 1] \\ 0.5 \cdot \frac{1}{(G_{\max}-1)} & G_T \in (1, G_{\max}] \\ 0 & \text{otherwise} \end{cases} \tag{6.16}$$

The posterior distributions of μ and σ are calculated using both the likelihood in Eq. 6.9 for the interpretation $G < 1$ and the likelihood in Eq. 6.11 for $G_{LB} < G < 1$. The quantiles of the posterior predictive distributions of M from both cases are compared with the true quantiles. Two cases are considered: one in which the analyst has estimated G_{LB} correctly, with $G_{LB} = 0.8$, and a second in which the analyst makes the same esti-

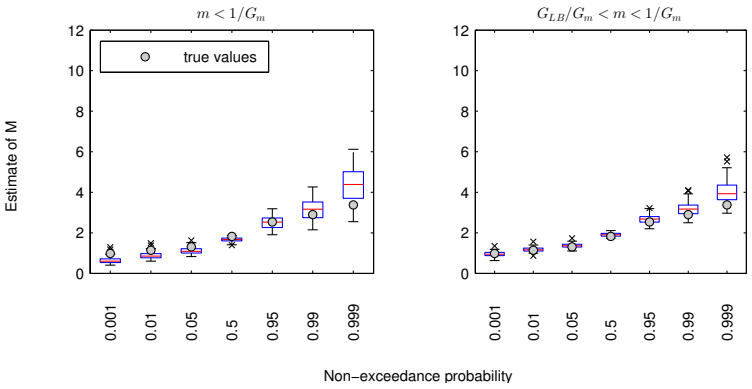

Figure 6.11: Quantiles of the model uncertainty M estimated for two interpretations of failure, for the case that $G_{LB} = 0.8$ is an accurate estimate. Both cases were run using a generated data set with 25 observations, and informed priors (prior case 2 described in Section 6.3.2).

mate, but incorrectly, where the true lower bound is $G_{LB} = 0.5$. For both cases, I use a generated data set of length $n = 25$, and use the informed priors (see Case 2 in Section 6.3.2).

The results are presented in Figs. 6.11 and 6.12. Figure 6.11 shows the case that G_{LB} has been correctly estimated. The left-hand plot shows the results when the likelihood for the constrained equality is used. The right-hand plot shows the results when the likelihood for the one-sided inequality interpretation of failure is used. In this case, where the estimate of G_{LB} is good, the quantiles of the posterior predictive distribution of M are more accurate, and show less variance, when using the constrained inequality interpretation of failure. However, if the estimate of G_{LB} is not good, the conclusion is the opposite. Figure 6.12 shows the case when the lower bound is estimated as $G_{LB} = 0.8$ but is actually 0.5. In this case, the constrained inequality interpretation of failure actually results in worse estimates of the quantiles of the M distribution. This is logical, because the constrained inequality interpretation only gives weight to samples of μ and σ when there is substantial probability mass of M within the interval $[G_{LB}/G_m, 1/G_m]$. When G_{LB} is overestimated, that interval is underestimated. Specifically, the left-hand side of Figure 6.12 shows that the quantile estimates of M are too high. This is because μ-σ samples that led to lower values of M were not given proper weighting due to the too-high estimate of G_{LB}. The take-home message is that when G_{LB} is well estimated, the constrained inequality improves the posterior predictive distribution of M. However, if estimated incorrectly, it can make the estimate worse. The use of a constrained inequality interpretation of failure should therefore be made exclusively on a well-founded basis.

6.4. APPLICATION

The BN method was used to estimate the uncertainty in the slope instability model D-Geostability ([78]). This application is fairly relevant because the model uncertainty is needed in the development of assessment tools for the primary flood defenses in the Netherlands. I obtained a data set of eleven locations where substantial site investigation

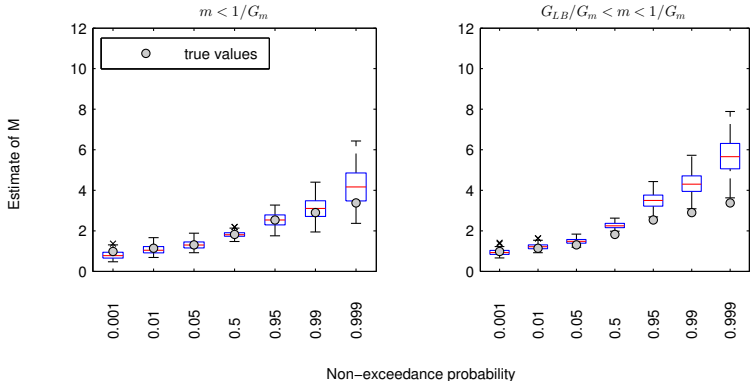

Figure 6.12: Quantiles of the model uncertainty M estimated for two interpretations of failure, for the case that $G_{LB} = 0.8$ is an incorrect estimate, with the true lower bound being 0.5. Both cases were run using a generated data set with 25 observations, and informed priors (prior case 2 described in Section 6.3.2).

had occurred, and where the loading data (either water levels or pore water pressures) were also well characterized for a particular performance observation (either failure or survival). For each location, a stability factor was computed with the slope instability model using the well characterized input data. In this method, it is assumed that all uncertainty in the stability factor is due to error in the model, and not to uncertainty in the input. In this application, a Lognormally distributed multiplicative model uncertainty M is used.

6.4.1. SLOPE STABILITY MODEL

Macrostability, or slope stability, is the resistance of the soil against shearing, which usually occurs along curved sliding planes, as schematized in Figure 6.13. The slope stability computations carried out in this chapter are based on Limit Equilibrium Modeling (LEM) and implemented in the D-Geostability model ([78]), specifically using the UpliftVan LEM. The soil can be modeled as drained or undrained. Upon compression, the former does not experience built up soil pore water pressures because it is assumed to drain indefinitely, which is typically representative of sandy soils. The latter, by contrast, does allow for built up pore pressures upon compression. In the D-Geostability model used in this research, undrained shear strength modeling based on Critical State Soil Mechanics is used to model the shear strength of peat and clay ([79]), specifically with the SHANSEP implementation (Stress History and Normalized Soil Engineering Properties). The regular drained shear strength modeling based on the friction angle is used for sand. The undrained shear strength depends on the following parameters in this model: (1) undrained shear strength ratio, (2) strength increase exponent and (3) yield stress, all of which can be determined using lab tests. Further details about the slope stability computation and parameters are beyond the scope of this chapter.

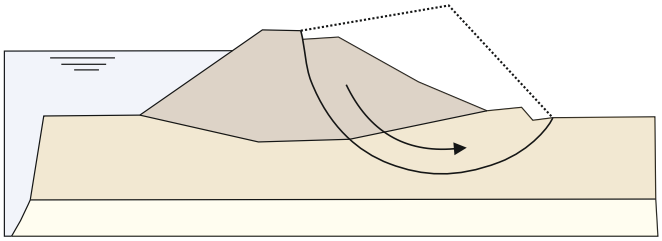

Figure 6.13: Slope instability on the land-side of the levee

6.4.2. DATA

The data set, which includes modeled stability factors and failure or survival observations, are provided in Table 6.2. All the locations are in the Netherlands; the dates range from 1928 to 2015. Most of the data are from historical events, but three are from field experiments. In the historical cases, the water levels during the storm event were available, but not the exact moment of failure. The values in Table 6.2 are for the peak load during the storm. In this section a few important details about the data are specified which are relevant for the model uncertainty analysis.

WOLPHERENSEDIJK 1980

The levee at Wolpherensedijk failed in 1980, but the soil measurements were taken after the levee had been restored. To compensate, the yield stress was lowered in the slope stability calculation. However, it is unclear if this compensation is valid, and it is therefore noted here that this data point brings uncertainty with it.

BERGAMBACHT 2001

At Bergambacht, a field test was carried out in 2001, first in September, and then again in November. These experiments are referred to as BA1 and BA2, respectively. In the BA1 experiment, the load required to fail the levee was underestimated and the loading capacity proved to be insufficient. Hence, the experiment was redone (BA2) in November, with a higher loading capacity. Furthermore, two cross sections were analyzed from the field test, a western one and an eastern one. Because these two locations are correlated, only one is included in the analysis. This is because the likelihood described by Eq. 6.8 assumes independence between failure locations.

For the BA1 experiment, the stability factor for the eastern cross section was G_m = 0.76, which is quite low considering there was no failure. The stability factor for the western cross section for the same experiment was G_m = 1.11. It is suspected that the strength of the western portion of the levee contributed to the survival of the eastern portion, due to 3D effects not accounted for by the D-Geostability model. This error is due to a shortcoming of the model, which is what the model uncertainty distribution is intended to account for. For this reason, even though the error in the BA1 computed stability factor is a bit of an outlier compared to the rest of the data in Table 6.2, it is included in the analysis. However, there has been some controversy about this data point, and some discussion about whether the calculation truly represented the situation during the field experiment. I felt that uncertainty about the validity of the BA1 data point

Figure 6.14: The macrostability field test on the IJkdijk in 2008.

warranted a separate analysis where this data point is excluded. Both results with and without the BA1 data point are presented in this section.

IJKDIJK 2008

The IJkdijk is an artificial levee that was built for testing purposes ([80]), and is only about 100 meters long. It therefore has edge effects that would not be present in a real levee. This might explain why the modeled stability factor for this field test was so low, $G_m = 0.64$, at the time of failure, when presumably $G = 1$. In this case, the error is not due entirely to the model, but to the test site not meeting the assumptions of the model (specifically: the assumption that the levee is long enough that the edges do not contribute to the stability).

6.4.3. SENSITIVITY TO DATA

Most of the modeled stability factors in Table 6.2 are in agreement with the failure or survival observations. However, as mentioned, one notable exception is the BA1 field experiment. In that case, failure did not occur (recall the load capacity was insufficient), but the calculated stability factor was $G_m = 0.76$. In this case of levee survival, $G > 1$ and $M > 1/G_m$, which in this case is $M > 1.3$. Inclusion of this data point results in a much wider variance in the distribution of M than its exclusion. Results are presented for the analysis both with and without the inclusion of the BA1 data point. It is advisable that those familiar with the experiment confirm that there were no peculiarities about the experiment that contributed to the suspiciously low value. If this can not be confirmed, then this data point should be included. This was a location where 3D effects, which are not accounted for in the D-Geostability model, may have had a large effect. In this case,

the large error for this experiment would genuinely be due to model uncertainty.

Table 6.2: Data set used in the BN application. Given are the locations, the location abbreviations (Abbr.), when relevant, the year of the observations, the calculated stability factors G_m, the observations, and their type (either historic or field test).

Location	Abbr.	Year	G_m	Observation	Type
Markermeerdijk	–	1928	1.00	Survived	Historic
Lekdijk Nieuw-Lekkerland	–	1953	1.00	Survived	Historic
Wolpherensedijk	–	1980	0.99	Failed	Historic
Streefkerk	–	1984	1.02	Failed	Historic
Lekdijk West	–	1995	1.00	Survived	Historic
Wolpherensedijk	–	1995	1.86	Survived	Historic
Bergambacht	BA0	1995	1.27	Survived	Historic
Bergambacht (Sep)	BA1	2001	0.76	Survived	Field test
Bergambacht (Nov)	BA2	2001	0.30	Failed	Field test
Zuiderlingedijk	–	2006	0.90	Failed	Historic
IJkdijk	–	2008	0.64	Failed	Field test

6.4.4. FAILURE INTERPRETATION

Section 6.2.2 discussed the interpretation of failure, in terms of what it means for the likelihood function. In the historic cases in this application, the peak load was used to derive the stability factors in Table 6.2. However, the load at the time of failure in these cases is unknown. Unless failure happened at the peak load during the storm, the computed stability factor at the actual failure load would be higher than the value in Table 6.2. Thus, in the case of historic failure, it is only known that $G < 1$, which means that $M < 1/G_m$.

For the field tests, a stronger interpretation of failure may be considered, such as $G = 1$ or $G_{LB} < G < 1$, where G_{LB} is the (expert-elicited) lower bound on what the true stability factor could be. All three approaches were carried out: $G < 1$, $G_{LB} < G < 1$, and $G = 1$, for the two data points from field experiments with failure observations, and the results were compared using cross validation (see Section 6.4.6 below), to see which interpretation leads to the best posterior predictive distribution of M. For the lower bound, $G_{LB} = 0.85$ for the field experiment BA2 (the November experiment at Bergambacht; see Table 6.2), and $G_{LB} = 0.7$ for the IJkdijk experiment. The BA2 bound is higher to reflect that there is no obvious reason to assume a deviation from $G = 1$, whereas at IJkdijk, there is. As discussed in the data section, the IJkdijk experimental levee is much shorter than real levees, while the D-Geostability model assumes an infinite length. Therefore, the experiment itself is responsible for some of the error, which reflects in the lower value chosen for G_{LB}.

6.4.5. PRIOR DISTRIBUTIONS

For the prior distributions of μ and σ (i.e. the parameters of the lognormal M distribution), I chose the four cases explored in the synthetic example in Section 6.3.2, with one change. The uniform distribution of σ was $\sigma \sim \mathcal{U}(0, 0.5)$ in the synthetic example; it is modified here to have a larger range: $\sigma \sim \mathcal{U}(0, 1)$. The parameters of the lognormal

distribution are not particularly intuitive, because they are the mean and standard deviation of $\ln(M)$, not M. To give more intuition about the selected priors, Figures 6.15 and 6.16 show the priors of μ and σ translated into the prior distributions of the mean and standard deviation of M, for the case of uniform priors and informed priors, repsectively. The relationships between the mean and standard deviation of M and the parameters μ and σ were given in Eqs. 6.14 and 6.15.

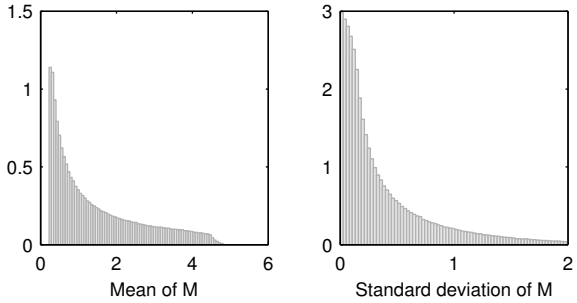

Figure 6.15: Prior distribution of the mean and standard deviation of the model uncertainty M, for prior case (1) where $\mu \sim \mathcal{U}(-1.5, 1.5)$ and $\sigma \sim \mathcal{U}(0, 1)$, and no expert constraint.

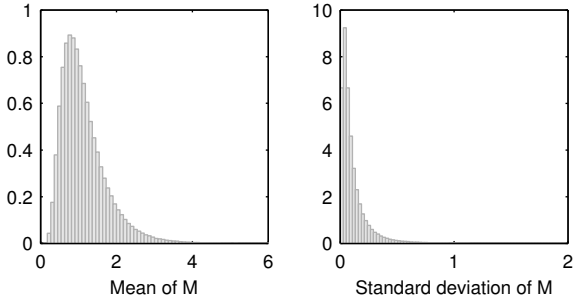

Figure 6.16: Prior distribution of the mean and standard deviation of the model uncertainty M, for prior case (1) where $\mu \sim \mathcal{N}(0, 0.5^2)$, $\sigma \sim \ln\mathcal{N}(0.1, 0.1^2)$, and no expert constraint.

6.4.6. CROSS VALIDATION

In total, 36 combinations were considered: three alternative interpretations of failure at two field test sites, and four alternative prior distributions. To determine which of these combinations has the best predictive capability, a cross validation was performed. This is a technique that is particularly useful when working with small amounts of data. Cross validation separates the dataset into a training set, and a validation set. In this application, I used a form called Leave-one-out cross validation, in which all locations except for one are used to determine F_m (the posterior predictive distribution of M), which is then validated using the observation at the left-out location (known as the validation location). The process is carried out iteratively, with each location serving as a validation

location one time. The score that each combination k receives at a validation location i is equal to $F_{m_k}(1/G_{m,i})$ (which is the probability of failure $P(G_i < 1)$) when failure was observed at the validation location, and equal to $1 - F_{m_k}(1/G_{m,i})$ when survival was observed. The total score of each combination is taken as the product over the scores at all of the validation locations. Eqs 6.17 and 6.18 summarize this information, where $s_{k,i}$ is the score of the k^{th} combination at the i^{th} validation location. F_{m_k} in Eq. 6.17 is the posterior predictive distribution of M calculated by the k^{th} combination, with the data point at the i^{th} location excluded.

$$s_{k,i} = \begin{cases} F_{m_k}\left(1/G_{m,i}\right), & \text{Fail}_\text{i} = 1 \\ 1 - F_{m,k}\left(1/G_{m,i}\right), & \text{Fail}_\text{i} = 0 \end{cases} \tag{6.17}$$

$$Score_k = \prod_i s_{k,i} \tag{6.18}$$

The cross validation score was calculated for each of the 36 combinations, and then converted to a relative score, in which the validation score is divided by the maximum validation score over all the combinations. This helps to easily compare the different combinations. The highest score will be 1, and the score of the other combinations indicates their goodness relative to the best combination. For example, a score of 0.8 means the combination was 20% worse than the best combination.

6

6.4.7. POSTERIOR RESULTS

In this section, I present the quantiles of the posterior predictive distribution of M, given the data and observations in Table 6.2. The posterior predictive distribution of M is calculated using numerical integration, according to Eq. 6.13. Tables 6.5 and 6.6 present the results for the case with the BA1 data point included and excluded, respectively. The tables present the specifics of each combination (prior type and failure interpretation at the two field experiment sites), as well as the 0.1%, 1%, 5%, 50%, 95%, 99%, and 99.9% quantiles (denoted, e.g., $Q_{0.1}$ for the 0.1% quantile), and the relative cross validation score. The combination(s) with the highest score are highlighted in the tables.

The best combination in the case where the BA1 data point is included is given in Table 6.3. In the case where the BA1 data point is excluded, there are two combinations which scored equally high in the cross validation. These are both given in Table 6.4. Posterior predictive densities of M for all three combinations given in Tables 6.5 and 6.6 are presented in Figure 6.17. The density when the BA1 data point is included is wider than the densities where it is excluded. Despite these differences, the 5% quantile in all three cases is fairly similar (0.7 when BA1 is included and 0.8 when it is excluded). Recall that this quantile is relevant in the Dutch safety assessments.

Table 6.3: Best performing combination of prior distributions and failure interpretation, for the case when the September Bergambacht data point (BA1) is included in the analysis.

Prior	Failure interpretation, Ijkdijk	Failure interpretation, Bergambacht
Uniform constrained	$M = 1/G_m$	$M < 1/G_m$

Table 6.4: Best performing combinations of prior distributions and failure interpretation, for the case when the September Bergambacht data point (BA1) is excluded in the analysis.

Prior	Failure interpretation, Ijkdijk	Failure interpretation, Bergambacht
Informed constrained	$M < 1/G_m$	$M < 1/G_m$
Informed constrained	$0.7/G_m < M < 1/G_m$	$M < 1/G_m$

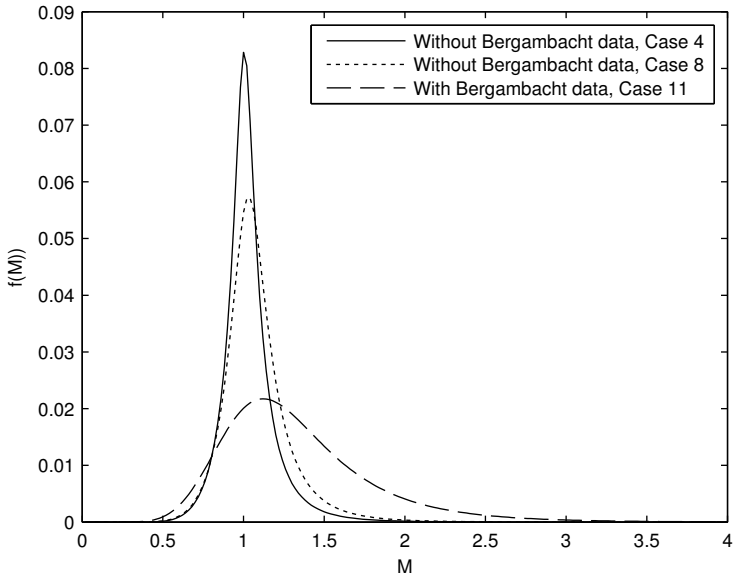

Figure 6.17: Posterior predictive distribution of the model uncertainty M for the best performing combinations of prior and failure interpretation, see Tables 6.5 and 6.6 for details about the cases.

With the inclusion of the BA1 data point, which was a survival observation and a modeled stability factor of 0.76, the best performing combination includes the the failure interpretation $M = 1/G_m$, which is understandable because this serves to shift the distribution to the right (higher values of M). The final cross-validation score is the product over the scores at the individual locations (see Eq. 6.18). Thus, the optimal combination in the case where BA1 is included is one that validates well at the locations besides BA1, but is still wide enough, or shifted to the right enough, to garner a decent score at BA1. This can be seen in the optimal choice (case 11, see Table 6.5). The variance is greater, and the upper tail much heavier than the optimal choices when BA1 is excluded (see Figure 6.17).

When the BA1 data point is excluded, there are two high-scoring combinations. One is where failure is interpreted as $M < 1/G_m$ at both the field experiment sites. The other is where failure at the IJkdijk is interpreted as $0.7/G_m < M < 1/G_m$, and failure for the BA2 case is interpreted as $M < 1/G_m$. Figure 6.17 shows that the difference between the posterior predictive distribution of M computed by these two is fairly minor.

Table 6.5: Quantiles of the posterior predictive distribution of the model uncertainty M for the 36 combinations of prior distributions, failure interpretation at IJkdijk (IJk.) and at Bergambacht (Berg.), for the case where the September survival data point at Bergambacht was included in the analysis. Q_x represents the x% quantile of M. The failure interpretations are (1) $F_1 : m < 1/Gm$, (2) $F_2 : G_{LB}/G_m < m < 1/G_m$, (3) $F_3 : m = 1/G_m$.

Case	Prior case	IJk.	Berg.	$Q_{0.1}$	Q_1	Q_5	Q_{50}	Q_{95}	Q_{99}	$Q_{99.9}$	Rel. Score
1	1	F_1	F_1	0.1	0.2	0.4	1.2	4.1	8.0	18.1	0.7
2	2	F_1	F_1	0.2	0.4	0.7	1.1	1.9	2.9	5.8	0.3
3	3	F_1	F_1	0.5	0.6	0.7	1.2	2.0	2.7	3.8	0.7
4	4	F_1	F_1	0.5	0.6	0.8	1.1	1.7	2.2	3.1	0.3
5	1	F_2	F_1	0.1	0.2	0.4	1.2	4.0	8.0	18.1	0.8
6	2	F_2	F_1	0.3	0.5	0.7	1.1	1.9	2.8	5.3	0.5
7	3	F_2	F_1	0.5	0.6	0.7	1.2	2.1	2.7	3.8	0.8
8	4	F_2	F_1	0.5	0.6	0.8	1.1	1.8	2.3	3.1	0.4
9	1	F_3	F_1	0.1	0.2	0.4	1.3	4.4	8.6	19.3	0.8
10	2	F_3	F_1	0.2	0.4	0.7	1.2	2.2	3.4	6.5	0.8
11	3	F_3	F_1	0.4	0.6	0.7	1.2	2.2	2.9	4.0	1.0
12	4	F_3	F_1	0.5	0.6	0.8	1.2	2.0	2.5	3.5	0.8
13	1	F_1	F_2	0.1	0.2	0.4	1.4	5.6	11.0	>20	0.7
14	2	F_1	F_2	0.1	0.3	0.5	1.3	3.5	6.0	12.4	0.9
15	3	F_1	F_2	0.4	0.6	0.7	1.4	2.6	3.4	4.7	0.9
16	4	F_1	F_2	0.4	0.6	0.7	1.3	2.5	3.2	4.4	0.9
17	1	F_2	F_2	0.1	0.2	0.4	1.5	5.7	11.1	>20	0.7
18	2	F_2	F_2	0.2	0.3	0.5	1.3	3.4	5.7	11.6	0.8
19	3	F_2	F_2	0.4	0.6	0.7	1.4	2.6	3.4	4.7	0.8
20	4	F_2	F_2	0.4	0.6	0.8	1.4	2.5	3.3	4.4	0.8
21	1	F_3	F_2	0.1	0.2	0.4	1.5	5.9	11.4	>20	0.6
22	2	F_3	F_2	0.2	0.3	0.5	1.4	3.5	5.9	11.9	0.8
23	3	F_3	F_2	0.4	0.6	0.8	1.4	2.6	3.5	4.7	0.7
24	4	F_3	F_2	0.4	0.6	0.8	1.4	2.5	3.3	4.5	0.7
25	1	F_1	F_3	0.1	0.2	0.4	1.4	5.8	11.4	>20	0.7
26	2	F_1	F_3	0.1	0.3	0.5	1.3	3.7	6.4	13.3	0.9
27	3	F_1	F_3	0.4	0.6	0.7	1.4	2.6	3.5	4.7	0.9
28	4	F_1	F_3	0.4	0.6	0.7	1.4	2.5	3.3	4.5	0.9
29	1	F_2	F_3	0.1	0.2	0.4	1.5	5.9	11.6	>20	0.7
30	2	F_2	F_3	0.2	0.3	0.5	1.4	3.6	6.2	12.6	0.8
31	3	F_2	F_3	0.4	0.6	0.7	1.4	2.6	3.5	4.7	0.8
32	4	F_2	F_3	0.4	0.6	0.8	1.4	2.6	3.3	4.5	0.8
33	1	F_3	F_3	0.1	0.2	0.4	1.6	6.1	11.9	>20	0.6
34	2	F_3	F_3	0.2	0.3	0.5	1.4	3.7	6.3	12.8	0.8
35	3	F_3	F_3	0.4	0.6	0.8	1.4	2.7	3.5	4.8	0.7
36	4	F_3	F_3	0.4	0.6	0.8	1.4	2.6	3.4	4.6	0.7

Table 6.6: Quantiles of the posterior predictive distribution of the model uncertainty M for the 36 combinations of prior distributions, failure interpretation at IJkdijk (IJk.) and at Bergambacht (Berg.), for the case where the September survival data point at Bergambacht was excluded in the analysis. Q_x represents the x% quantile of M. The failure interpretations are (1) $F_1 : m < 1/Gm$, (2) $F_2 : G_{LB}/G_m < m < 1/G_m$, (3) $F_3 : m = 1/G_m$.

Case	Prior case	IJk.	Berg.	$Q_{0.1}$	Q_1	Q_5	Q_{50}	Q_{95}	Q_{99}	$Q_{99.9}$	Rel. Score
1	1	F_1	F_1	0.1	0.2	0.4	1.0	3.2	6.3	14.5	0.5
2	2	F_1	F_1	0.3	0.6	0.8	1.0	1.4	1.8	3.2	0.9
3	3	F_1	F_1	0.5	0.6	0.7	1.0	1.7	2.2	3.2	0.9
4	4	F_1	F_1	0.6	0.7	0.8	1.0	1.3	1.6	2.2	1.0
5	1	F_2	F_1	0.1	0.2	0.4	1.1	3.2	6.4	14.8	0.5
6	2	F_2	F_1	0.4	0.6	0.8	1.1	1.5	2.0	3.5	0.9
7	3	F_2	F_1	0.5	0.6	0.7	1.1	1.8	2.4	3.4	0.9
8	4	F_2	F_1	0.5	0.7	0.8	1.1	1.4	1.8	2.5	1.0
9	1	F_3	F_1	0.1	0.2	0.4	1.2	3.8	7.4	16.7	0.5
10	2	F_3	F_1	0.2	0.4	0.6	1.1	2.0	2.9	5.4	0.7
11	3	F_3	F_1	0.4	0.6	0.7	1.2	2.1	2.7	3.8	0.7
12	4	F_3	F_1	0.5	0.6	0.7	1.1	1.8	2.3	3.2	0.8
13	1	F_1	F_2	0.1	0.2	0.3	1.3	5.1	9.9	>20	0.3
14	2	F_1	F_2	0.1	0.3	0.5	1.2	3.2	5.6	11.5	0.5
15	3	F_1	F_2	0.4	0.6	0.7	1.3	2.5	3.4	4.6	0.4
16	4	F_1	F_2	0.4	0.6	0.7	1.3	2.4	3.2	4.3	0.4
17	1	F_2	F_2	0.1	0.2	0.4	1.4	5.2	10.2	>20	0.3
18	2	F_2	F_2	0.2	0.3	0.5	1.3	3.2	5.4	10.9	0.4
19	3	F_2	F_2	0.4	0.6	0.7	1.4	2.6	3.4	4.6	0.3
20	4	F_2	F_2	0.4	0.6	0.7	1.3	2.4	3.2	4.4	0.3
21	1	F_3	F_2	0.1	0.2	0.4	1.4	5.4	10.5	>20	0.3
22	2	F_3	F_2	0.2	0.3	0.5	1.3	3.3	5.6	11.2	0.4
23	3	F_3	F_2	0.4	0.6	0.7	1.4	2.6	3.4	4.7	0.3
24	4	F_3	F_2	0.4	0.6	0.7	1.4	2.5	3.2	4.4	0.3
25	1	F_1	F_3	0.1	0.2	0.3	1.3	5.3	10.3	>20	0.3
26	2	F_1	F_3	0.1	0.2	0.4	1.2	3.4	6.0	12.4	0.5
27	3	F_1	F_3	0.4	0.6	0.7	1.4	2.6	3.4	4.7	0.4
28	4	F_1	F_3	0.4	0.6	0.7	1.3	2.5	3.3	4.5	0.4
29	1	F_2	F_3	0.1	0.2	0.4	1.4	5.5	10.6	>20	0.3
30	2	F_2	F_3	0.1	0.3	0.5	1.3	3.4	5.8	11.8	0.4
31	3	F_2	F_3	0.4	0.6	0.7	1.4	2.6	3.4	4.7	0.3
32	4	F_2	F_3	0.4	0.6	0.7	1.3	2.5	3.3	4.5	0.3
33	1	F_3	F_3	0.1	0.2	0.4	1.4	5.6	10.9	>20	0.3
34	2	F_3	F_3	0.1	0.3	0.5	1.3	3.5	6.0	12.1	0.4
35	3	F_3	F_3	0.4	0.6	0.7	1.4	2.6	3.5	4.7	0.3
36	4	F_3	F_3	0.4	0.6	0.7	1.4	2.5	3.3	4.5	0.3

6.5. DISCUSSION AND CONCLUSIONS

In this chapter, a methodology was presented to use a BN to derive the model uncertainty in a failure model using observations and hindcasts. The choice of prior and the impact of limited data were investigated via a synthetic example. It was found that small data sets can result in large variance for high quantiles if the priors are poorly chosen. That is, the smaller the data set, the more important it becomes to choose a well-informed prior. However, for most quantiles, the posterior predicted estimates were quite accurate, even for relatively small data sets. When little is known apriori about the parameters of the model uncertainty, experts can impart their intuition in the form of minimum and maximum quantiles of M for specified exceedance probabilities. Incorporating this intuition delivers a lot of reduction in the variance of the estimated quantiles of M.

The issue of how to interpret failure observations for field experiments was addressed. In cases where the soil parameters are well known and the load at the time of failure is recorded, it is reasonable to assume that $G = 1$ has been observed. When there is uncertainty about any aspect of the experiment, this can become an overly confident assumption, so the interpretation that $G_{LB} < G < 1$ was also tested. This requires some intuition about what the lower bound G_{LB} should be. I showed in the synthetic example that if G_{LB} is estimated well, choosing this interpretation over the more general (but always correct) interpretation that $G < 1$ can reduce the variance in the posterior predicted quantiles of M, and improve their accuracy. However, if G_{LB} is incorrectly chosen, it can have a rather negative effect, resulting in poorer estimates and larger variance.

The method was applied to the slope stability model D-Geostability, using data at 11 locations in the Netherlands. One of these data points was potentially an outlier. It was during a field experiment, in which the levee did not fail, but the calculated stability factor was 0.76. Most of the other data is either on the correct side of 1 (i.e. $G_m < 1$ for failure or $G_m > 1$ for survival), or close to it. There is also a lot of skepticism about this modeled stability factor. Those familiar with the experiment were not in close contact with those doing the hindcast modeling, and it is unclear if the correct conditions were represented in the model. Because of this, the posterior results of M were calculated using the data both with and without the outlier data point. The results were quite sensitive to the inclusion or exclusion of this data point, which underscores the importance of using trustworthy data.

In general, the method presented in this chapter is fast and easy to implement. The synthetic example showed how with relatively little data, the method can well estimate the correct quantiles of the model uncertainty distribution. I recommend applying this method whenever quality data are available as input to the failure model, together with survival or failure observations.

7

CONCLUSIONS

In this dissertation, the primary interest is in the development and use of a Bayesian network (BN) for levee system reliability to augment and support existing methods. In particular, the research focused on the verification of key algorithms in the Hydra-Ring reliability model (or its predecessor PC-Ring), which is at the heart of the Dutch national flood risk analysis and flood defense assessments. The BN developed in this dissertation works with MC sampling, where the correlation structure is realized with a Gaussian copula. In this sense, the BN can be considered a more explicitly-modeled MC approach with which to compare Hydra-Ring reliability estimates. The benefit of representing the method in a BN are threefold: (1) The correlation structure, which can become complex for large systems, is more clear and explicit when using a network representation, (2) The network representation makes it possible for non-probabilistic-experts to understand the problem and be able to impart their own expertise (consider, for example, geotechnical experts, who may know a lot about the failure mechanism, but not understand the probabilistic techniques), and (3) the graphical representation allows one to determine both prior and posterior dependences/independences among the variables in the network, using a concept known as D-separation. Appendix D discusses this third point in detail. In addition to verifying algorithms in Hydra-Ring, the BN is able to improve reliability estimates at a system scale, which was done in both synthetic and real-world examples in the dissertation. Prior to this research, such updating had only been done at a cross-sectional scale. The research also focused on how the Hydra-Ring algorithms could be used in combination with survival observations (high water levels which a levee survived) to update the reliability at the system scale, and whether such methods are accurate. As with the BN, this was tested for both synthetic and real-world examples. The last research topic in the dissertation looked at the use of a BN to quantify uncertainty in a geotechnical failure model, using hindcasted model output together with historically observed failure or survival. Uncertainty in the failure model can dominate the failure probability; using a BN and observations to reduce or better quantify that uncertainty can therefore be strongly beneficial to uncertainty reduction in the failure probability estimate.

In the following section, the main findings from the research in the dissertation are presented. Thereafter, recommendations for further research are discussed.

7.1. MAIN FINDINGS

DEVELOPMENT OF A BN FOR LEVEE SYSTEM RELIABILITY

Chapter 2 was dedicated to describing the methodology to build a BN for a spatially-extensive levee system. The main challenge in levee system reliability modeling with a BN is the densely correlated variables in the network, which results from the spatial autocorrelation of the resistance variables. Densely correlated variables translates to lots of arcs in the network, and child nodes with many parents. The traditional and most common BN - the discrete BN - would contain conditional probability tables in these cases that could easily have trillions of conditional probabilities to specify, even for relative coarse discretization of the marginal distributions, making traditional networks an infeasible option for levee system reliability.

Chapter 2 provided background about different types of BNs, specifics about the BN used in this dissertation, and a detailed methodology for describing a spatially extensive levee system using the BN. The BN used in the dissertation was chosen because it allows the variables to be described by marginal continuous or discrete probability distributions, with dependency specified by correlation coefficients, and the dependence structure built using copulas. This type of network made the reliability modeling of spatial levee systems possible.

In the BN, the joint distribution is obtained, and inference is performed, using MC sampling. In fact, as discussed in Chapter 2, MC sampling with a Gaussian copula could be used without the BN representation. However, setting up the problem without the network would feel like a handicap, as it helps clarify the connections (correlations) within the joint distribution, both to the analyst and to the world in which he/she operates. This latter is essential in communication with experts in the field who are not experts in probabilistic analysis. Equally important is the ability to determine both prior and posterior independence/dependence among the variables in the network, which can be done by observing the graphical structure using a concept known as D-separation (see Appendix D). Particularly for the posterior case, after certain variables have been observed, determining the dependences without the graph would be cumbersome and prone to errors. For these reasons, I find the network representation of the joint distribution essential, and highly valuable. In the dissertation, the method is referred to as the BN method. However, when used to compare with Hydra-Ring algorithms, it can still be thought of as approximate efficient methods on the one hand (Hydra-Ring), and more explicit MC sampling on the other.

The following specific considerations were addressed in detail:

- Structuring/building the BN for different spatial scales: a cross section, a statistically homogeneous segment, and a levee system
- Sampling the network, accounting for correlations between variables, to derive samples of the multivariate (spatial) distribution over all the variables in the network

- Performing inference in the BN for different types of observations in the real world, like the survival of (high) loads
- The potential need for importance sampling in the estimation of the posterior failure probability
- Posterior dependence of the resistance variables

VERIFICATION OF THE MODIFIED OUTCROSSING METHOD

The BN methodology proposed in Chapter 2 was used to calculate the length effect in a levee segment. The length effect refers to the increase in failure probability that results from an increase in the length of a levee segment. It is caused by the spatial variability within a segment, and the probability of encountering a weak spot. The BN was used to calculate the length effect by sampling the joint spatial distribution of the limit state function, without having to approximate a parametric form of the spatial distribution. Using Monte Carlo rejection sampling for inference, the method was used to update failure probabilities of (long) levees using survival observations (i.e. high water levels and no levee failure). The BN results were compared with a modified outcrossing (MO) technique, currently in use in reliability modeling of flood defenses in the Netherlands, via a numerical example, for verification purposes. The primary difference between the two methods is that the BN method samples from the joint spatial distribution, whereas the MO method uses an approximative parametric form of the spatial distribution of the limit state, and solves the problem analytically.

The prior and posterior segment failure probabilities calculated by the two methods were in strong agreement. Slight discrepancies were found for posterior segment failure probabilities for long segments (4000 and 6000 meters), but these differences were less than 10%, and in terms of reliability index, less than 1%. These results provide a strong verification of the MO method for prior analysis, which is used in the levee reliability model Hydra-Ring that supports Dutch national flood risk assessments. They also provide an important verification of the MO method for posterior analysis, which has a lot of potential. The speed of the MO method makes it possible to efficiently update failure probabilities of numerous levee segments with abundant survival observations.

Given the strong agreement between BN and MO results, and the relative efficiency of the MO method, it is logical to use the latter in practice. However, I must emphasize that the examples considered in Chapter 3 do not represent an exhaustive set of cases. For failure probability updating with survival observations, I advocate comparing the BN and MO output for each new type of application (e.g. new limit state function, new set of variable distribution types or correlation parameters). Once the results are verified, the MO method can be used with confidence for all examples of the same type.

Finally, I strongly advocate the use of either the BN or MO method to account for the length effect in reliability analysis over some of the more simplified approaches found in the literature, which tend to break up the segment into fully correlated or completely independent sub-sections, such as the methods described in [45] and [46].

VERIFICATION OF THE EQUIVALENT PLANES METHOD

The Equivalent Planes (EP) method is an efficient method to compute the system reliability for systems of correlated components. It has been used in Dutch system reliability modeling of flood defenses for decades, in the model PC-Ring, and more recently Hydra-Ring. The latter is at the heart of national flood risk analysis, the results of which are used to drive major flood prevention policies in the Netherlands. The critical role of the model motivated this research to determine the accuracy of the EP method, and under which situations unacceptable error may be encountered. For numerical cases where the components were equi-correlated, exact solutions were available. In all other cases, Monte Carlo (MC) directional sampling method was used to compute 'exact' reliability estimates with which to compare the EP results.

The error in the EP method was computed for different system configurations, and found that when the components are not too correlated (e.g. a correlation coefficient up to about 0.5), the error in the method is generally negligible, particularly when the components have high reliability indices. Inaccuracies become apparent for large systems with highly correlated components, and for components with lower reliability indices. In all cases, the Equivalent Planes system failure probability estimates were within a factor of three times the correct system failure probability. It is important to note that these results are for extreme system configurations in which the components all have equal reliability indices. In reality, a few components will likely dominate the failure probability, and the error will be much lower. Furthermore, even three times the correct failure probability can be quite negligible for systems with very small failure probabilities. For example, consider a system of 250 equi-correlated components, with component reliability indices of 6, correlated with a coefficient of 0.9; the true failure probability is 4.81E-8 and the estimate is 1.14E-7, for a factor difference of 2.4. In many applications, where the probability needs to be below a certain safety standard, this difference will not be important. Furthermore, other uncertainties in the reliability analysis – for example, due to the parameterization of the random variables contributing to failure – will likely overshadow this small error, making it essentially negligible.

The EP method is particularly attractive because of its efficiency. For a system with 250 components, all with equal reliability indices of 5, and equi-correlated with correlation coefficient 0.9, the EP method computed the system failure probability in 0.95 seconds[1]. For the same case, MC directional sampling (which is already more efficient than crude MC) required 1079 minutes – over 17 hours. The other attractive feature of the EP method is that it returns influence coefficients, which makes it particularly unique among the (computationally feasible) methods developed within the field of system reliability (see Section 4.1 for references). While there are methods to return influence coefficients of the components, the EP method additionally returns influence coefficients of the random variables. This lets managers know which variables (and in which segments) are most contributing to the failure probability of the system.

A BN was also applied to the case of equi-correlated components. Because the limit state functions are all approximated as normal distributions in the examples in Chapter 4 (this is the first step of the EP method), the BN of a system of m components is equivalent to a multivariate normal distribution. At the time of the development of the EP method,

[1] Computation times are based on a 2.8 GHz computer with 8GB RAM

evaluation of the multivariate normal distribution for high reliability indices and a large number of components was not feasible. Since that time, there have been developments, specifically methods by Genz and Bretz ([70], [71]) which have been adopted in statistical packages, such as Matlab (which was used in this research). The method is known as a quasi-MC method of evaluating multivariate normal distributions. It performed well for reliability indices of 3 and 4, for all component correlations. For an index of 5, the method performs well until the correlation gets high (> 0.7), and for a reliabiity index of 6, the method only performs accurately when the components are barely correlated ($\rho = 0.2$). The poor performance appears to be due to an inability of the quasi-MC method to accurately calculate the error in the probability estimate in some cases. For example, for a system of 250 components, with component reliability indices of 6, and equi-correlated with coefficient 0.9, the exact system failure probability is about $5 \cdot 10^{-8}$, while the quasi-MC method returns an estimate of about $1 \cdot 10^{-9}$. However, the method estimates an error of about 10^{-12}, which is clearly incorrect. Details about the quasi-MC algorithm to estimate system probabilities and error fall outside the scope of this dissertation. Based on the results of our application of it to the equi-correlated case in Chapter 4, I conclude that directly evaluating the multivariate normal distribution is still too unstable for practical application. However, I also recommend keeping up with the literature, as new algorithms to evaluate multivariate normal distributions continue to be developed.

UPDATING PIPING PROBABILITIES IN LEVEE SYSTEMS

The BN methodology described in Chapter 2 was applied to two levee systems in The Netherlands. The first system was composed of regional levees protecting the city of Heerhugowaard from a canal. The prior failure probabilities due to piping were nearly 1, although no evidence of failure had been seen. Applying the BN to the three segments that made up this system, and imposing the survival observation (a high water level with no failure), reduced the failure probabilities by over two orders of magnitude. Discussions with the water board indicated that it was not necessarily poor prior distributions of the soil parameters, but rather an inappropriate choice of model (the Sellmeijer piping model) for the particular levee system. Specifically, they believe a mucky impermeable layer exists under the canal, which disconnects the pressure of the water level from the underlying sand layer. In this case, the prior distributions of the soil parameters may be correct, and possibly quite weak; however, the disconnect makes it so that the load is not really felt, and therefore no failure occurs. This highlights an important point when applying the BN methodology: pay attention to the assumptions of the failure model. It is critical to ensure that the location where it is applied meets those assumptions. Otherwise, the resulting posterior distributions of the soil parameters will be nonsense.

The primary levee system that was considered, which protects an area of Zutphen against the IJssel River, was a more valid system for use of the Sellmeijer piping model, as it is unlikely that the IJssel river is underlain by an impermeable layer. The challenge with the primary system is that the prior probabilities are already quite low, so that a survival observation is not guaranteed to have much of an impact on the failure probability. For this case, posterior segment and system failure probabilities were calculated using two load observations (coupled with levee survival): a 1/40 year water level and a 1/400 year

water level, to assess how the extremity of the observation influences the reduction in failure probability.

In both the regional and primary cases, prior and posterior cross-sectional, segment, and system failure probabilities were calculated. In the regional case, the system failure probability decreased by over two orders of magnitude, due to the extremely high prior failure probability of this system. In the primary case, the system failure probability decreased by 7.5 for a 1/400 year water level observation, and by 2 for a 1/40 year water level observation. While a 1/40 year water level observation is relatively high, it was not high enough to greatly impact a system with a very low prior failure probability.

For the primary case, a further verification of the Hydra-Ring MO and EP algorithms was done by comparing their segment and system failure probability estimates with those calculated with the BN. Chapter 3 had explored the validity of the MO method by investigating numerical examples. Chapter 5 provided further verification by comparing the results for this real-world application. Though the prior segment failure probabilities showed slightly worse agreement than the numerical examples in Chapter 3, they were still in good agreement. In terms of reliability index, the MO prior estimates for the three segments in the system were 2%, 1%, and 3% lower (note: lower reliability index = higher failure probability) than the BN estimates. The MO posterior estimates (for the case of a 1/400 year water level) were in even better agreement, a mere 0.3%, 1%, and 1% lower than the BN estimates. The agreement between system failure probability estimates calculated by the BN and the combined MO/EP methods (MO for the segment, and EP for the system) was also investigated. In terms of reliability index, the differences were only about 1%. This difference is on par with the differences between the MO and BN segment estimates. It was expected that the EP method would be fairly exact for a system of only three components, based on the results in Chapter 4, so the good agreement is not surprising. Still, it is a useful verification in an actual application, as opposed to the numerical examples explored in Chapter 4.

USEFULNESS OF SURVIVAL OBSERVATIONS

Survival observations are vastly available, but not always particularly useful. In Chapter 5, the usefulness of survival observations under different conditions was investigated. Two primary factors determine how useful a survival observation will be: (1) the magnitude of the prior conditional failure probability given the observed load, and (2) the relative influence that the resistance has on failure. When the prior conditional failure probability is high, then observing survival is often very informative, because it tells us that our prior distributions are underestimating the strength of the levee. On the other hand, if the expected failure probability given the observed load is low, then observing survival essentially confirms our prior knowledge, and the posterior distributions are similar to the priors. For the second factor, if the resistance has little influence on the failure probability (i.e. the load is dominant), then improving its distribution will have little effect. Conversely, if the resistance is dominating the failure probability, then improving its distribution should have a large impact. The results for an analysis of a standard limit state function were presented in Section 5.7, but are repeated here, because they can serve as a useful guidance for anyone considering updating levee failure probability with survival observations. Consulting these tables prior to carrying out reliability

updating can help avoid unnecessary calculations.

Table 7.1: Reduction in failure probability (ratio of prior to posterior) due to incorporation of a survived load with a return period of **400 years**, for different influences of the resistance (10%, 30%, 70%, and 90%), and different prior reliability indexes β.

	$\alpha_R^2 = 0.9$	$\alpha_R^2 = 0.7$	$\alpha_R^2 = 0.5$	$\alpha_R^2 = 0.3$	$\alpha_R^2 = 0.1$
$\beta = 2$	2656	216	41	15	10
$\beta = 3$	1460	67	8	2	1
$\beta = 4$	839	23	2	1	1

Table 7.2: Reduction in failure probability (ratio of prior to posterior) due to incorporation of a survived load with a return period of **40 years**, for different influences of the resistance (10%, 30%, 70%, and 90%), and different prior reliability indexes β.

	$\alpha_R^2 = 0.9$	$\alpha_R^2 = 0.7$	$\alpha_R^2 = 0.5$	$\alpha_R^2 = 0.3$	$\alpha_R^2 = 0.1$
$\beta = 2$	245	26	6	2	1
$\beta = 3$	145	11	2	1	1
$\beta = 4$	91	5	1	1	1

BNS FOR ESTIMATING GEOTECHNICAL FAILURE MODEL UNCERTAINTY

A BN was developed to derive the model uncertainty in a geotechnical failure model using observations and hindcasts. The research investigated the choice of prior distributions for the distribution parameters of the model uncertainty, and the impact of limited data, via a synthetic example. Small data sets can result in large variance for high posterior predictive quantiles of M if the prior distributions are poorly chosen. That is, the smaller the data set, the more important it becomes to choose a well-informed prior. However, for most quantiles, the estimates were quite accurate, even for relatively small data sets. Even when little is known apriori about the parameters of the model uncertainty distribution, experts can impart their intuition in the form of minimum and maximum quantiles of M for specified exceedance probabilities. Incorporating this intuition delivers a lot of reduction in the variance of the posterior predictive quantiles of M.

The methodology was applied to a case study in the Netherlands using the marcro-instability model D-Geostability. The results showed a sensitivity to outliers, and underscored the need to use high-quality well-vetted data points when doing the analysis.

In general, the method is fast and easy to implement. The synthetic example showed how with relatively little data, the method can well estimate the correct quantiles of the model uncertainty distribution. I consider it worthwhile to apply this method whenever quality data are available as input to the failure model, together with survival or failure observations.

7.2. RECOMMENDATIONS FOR FUTURE RESEARCH

The BN developed in this dissertation has served as a yardstick by which to measure the goodness of more approximative, but potentially faster methods. I hope that it continues to serve in this fashion, and recommend it as an tool to augment any study considering posterior analysis of levee failure probabilities based on observations. Due to time constraints, certain issues were excluded from this dissertation, but their investigation would be useful in the future. These are outlined below.

PARAMETERIZATION OF POSTERIOR DEPENDENCE

The posterior dependence structure of the resistance variables is highly constrained and has no obvious parametric form. For that reason, in this dissertation, the posterior joint samples of the resistance variables (which were retained after rejection sampling) were used to calculate the posterior failure probability. This limits the accuracy of the estimate, because the number of retained samples can be quite low in some cases. I feel convinced that a method can be derived to parameterize the constrained posterior dependence that will allow for sampling the joint posterior distribution, which will allow us to sample as many times as needed.

APPLICATION WITH AN EXTERNAL FAILURE MODEL

In this dissertation, the real-world applications considered the piping failure mechanism, which is represented by empirical formulas. The BN developed for levee reliability in this dissertation is specifically described for such a case, with a formulaic representation of the limit state function. In Chapter 2, I described how the method can be expanded in the case that an external failure model (e.g. a macrostability model) is needed. It would require a preprocessing step in which the failure model is run with random input to derive the empirical distribution of the limit state function and the correlation between the input random variables and the limit state function. The limit state function would then enter the BN not as a functional node, but as a random node specified by its empirical distribution. Arcs would be drawn between the limit state function and the input variables, and specified by the correlation derived in the preprocessing step. From there, the method described in the dissertation could be applied.

The one caveat is that there would be correlations between variables within a network in addition to the autocorrelations that densely connect the cross sections in the BN. This would have to be carefully accounted for so that all correlations are adhered to when sampling, and so that the correlation matrix over all the variables in the system remains positive definite. Although such a situation did not arise in the applications in this dissertation, a method for deriving a positive definite matrix was described and tested in Appendix G.

TIME-DEPENDENT FAILURE MECHANISMS

The piping mechanism, which was considered in the BN applications in this dissertation, is not particularly sensitive to antecedent conditions. For a mechanism like slope stability which depends on how saturated the soil already is at the time of a load, a survival observation has a different meaning. Specifically, surviving a particular load does

not guarantee survival of the same load in the future, as it did in the case of piping. It will depend on the starting saturation state of the soil. Future research should focus on how inference would need to be adjusted in such a case.

7.3. FINAL THOUGHTS

It is my hope that this dissertation contributes in some small part to the current focus on incorporating survival observations in levee reliability estimates. Uncertainty is one of our biggest challenges in reliability analysis. And these survival observations are sitting there, waiting to help. I also hope that this dissertation both clarifies and increases the confidence in the modified outcrossing and Equivalent Planes algorithms that are a part of the flood defense reliability model Hydra-Ring.

REFERENCES

REFERENCES

[1] Heiko Apela, Annegret H Thiekena, Bruno Merza, and Günter Blöschlb. A probabilistic modelling concept for the quantification of flood risks and associated uncertainties. In *International Congress on Environmental Modelling and Software*, page 107, 2004.

[2] Heiko Apel, Bruno Merz, and Annegret H Thieken. Quantification of uncertainties in flood risk assessments. *International Journal of River Basin Management*, 6(2): 149–162, 2008.

[3] Bruno Merz and Annegret H Thieken. Flood risk curves and uncertainty bounds. *Natural hazards*, 51(3):437–458, 2009.

[4] Mary W Downton, Rebecca E Morss, Olga V Wilhelmi, Eve Gruntfest, and Melissa L Higgins. Interactions between scientific uncertainty and flood management decisions: Two case studies in colorado. *Global Environmental Change Part B: Environmental Hazards*, 6(3):134–146, 2005.

[5] Jim W Hall, RJ Dawson, PB Sayers, C Rosu, JB Chatterton, and R Deakin. A methodology for national-scale flood risk assessment. *Proceedings of the ICE-Water and Maritime Engineering*, 156(3):235–247, 2003.

[6] RB Jongejan and B Maaskant. Quantifying flood risks in the netherlands. *Risk Analysis*, 35(2):252–264, 2015.

[7] WT Bakker and JK Vrijling. Probabilistic design of sea defences. *Coastal Engineering Proceedings*, 1(17), 1980.

[8] HMGM Steenbergen, BLb Lassing, ACWM Vrouwenvelder, and PH Waarts. Reliability analysis of flood defence systems. *Heron, 49 (1)*, 2004.

[9] BL Lassing, ACWM Vrouwenvelder, and PH Waarts. Reliability analysis of flood defence systems in the netherlands. In *van Gelder PHAJM, Bedford T.(eds.), Safety and Reliability, Proc. Of the ESREL 2003 conference*, pages 1005–1014, 2003.

[10] Ton Vrouwenvelder. Spatial effects in reliability analysis of flood protection systems. In *Second IFED Forum, Lake Louise, Canada*, 2006.

[11] FLM Diermanse, W van Balen, K Roscoe, J Lopez de la Cruz, and H Steenbergen. Hydra-Ring 2.0, Probabilistics toolbox for the WTI2017, Scientific Document. Technical report, Deltares, 2016.

[12] J Ferry Borges, M Castanheta, and AR Janeiro Borges. *Design criteria for wind loads on statistical Bases*. 1971.

[13] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers, San Francisco, Calif., 1997. ISBN 1558604790 9781558604797.

[14] Irad Ben-Gal. Bayesian networks. In *Encyclopedia of Statistics in Quality and Reliability*. John Wiley & Sons, Ltd, 2008. ISBN 9780470061572. URL http://onlinelibrary.wiley.com/doi/10.1002/9780470061572.eqr089/abstract.

[15] Rina Dechter. Bucket elimination: A unifying framework for reasoning. *Artificial Intelligence*, 113(1–2):41–85, September 1999. ISSN 0004-3702. doi: 10.1016/S0004-3702(99)00059-4. URL http://www.sciencedirect.com/science/article/pii/S0004370299000594.

[16] Zhaoyu Li and Bruce D'Ambrosio. Efficient inference in bayes networks as a combinatorial optimization problem. *International Journal of Approximate Reasoning*, 11 (1):55–81, July 1994. URL http://www.sciencedirect.com/science/article/pii/0888613X94900191.

[17] N. L. Zhang and D. Poole. Exploiting causal independence in bayesian network inference. *Journal of Artificial Intelligence Research*, 5:301–328, November 1996. URL http://arxiv.org/abs/cs/9612101.

[18] A. Christofides, B. Tanyi, S. Christofides, D. Whobrey, and N. Christofides. The optimal discretization of probability density functions. *Computational Statistics & Data Analysis*, 31(4):475–486, October 1999. URL http://www.sciencedirect.com/science/article/pii/S0167947399000432.

[19] Nir Friedman and Moises Goldszmidt. Discretizing continuous attributes while learning bayesian networks. In *In Proc. ICML*, page 157–165. Morgan Kaufmann, 1996.

[20] Alexander V Kozlov and Daphne Koller. Nonuniform dynamic discretization in hybrid networks. In *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence*, pages 314–325. Morgan Kaufmann Publishers Inc., 1997.

[21] Helge Langseth, Thomas D Nielsen, Rafael Rumí, and Antonio Salmerón. Inference in hybrid bayesian networks. *Reliability Engineering & System Safety*, 94(10):1499–1509, 2009.

[22] Barry R. Cobb, Prakash P. Shenoy, and Rafael Rumí. Approximating probability density functions in hybrid bayesian networks with mixtures of truncated exponentials. *Statistics and Computing*, 16(3):293–308, September 2006. URL http://link.springer.com/article/10.1007/s11222-006-8175-8.

[23] Serafín Moral, Rafael Rumi, and Antonio Salmerón. Mixtures of truncated exponentials in hybrid bayesian networks. In Salem Benferhat and Philippe Besnard, editors, *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, number 2143 in Lecture Notes in Computer Science, pages 156–167. Springer Berlin Heidelberg, January 2001. ISBN 978-3-540-42464-2, 978-3-540-44652-1. URL http://link.springer.com/chapter/10.1007/3-540-44652-4_15.

[24] Rafael Rumí, Antonio Salmerón, and Serafín Moral. Estimating mixtures of truncated exponentials in hybrid bayesian networks. *Test*, 15(2):397–421, September 2006. URL http://link.springer.com/article/10.1007/BF02607059.

[25] Rafael Rumí and Antonio Salmerón. Approximate probability propagation with mixtures of truncated exponentials. *International Journal of Approximate Reasoning*, 45(2):191–210, July 2007. URL http://www.sciencedirect.com/science/article/pii/S0888613X06000594.

[26] Anca Hanea, Oswaldo Morales Napoles, and Dan Ababei. Non-parametric bayesian networks: Improving theory and reviewing applications. *Reliability Engineering & System Safety*, 144:265–284, 2015.

[27] Anca Maria Hanea. *Algorithms for non-parametric Bayesian belief nets*. PhD thesis, TU Delft, Delft University of Technology, 2008.

[28] D. Kurowicka and R.M. Cooke. Distribution-free continuous bayesian belief nets. *Modern Statistical and Mathematical Methods in Reliability*, 10:309, 2005.

[29] A.M. Hanea, D. Kurowicka, and R.M. Cooke. Hybrid Method for Quantifying and Analyzing Bayesian Belief Nets. *Quality and Reliability Engineering International*, 22(6):613–729, 2006.

[30] A.M. Hanea, D. Kurowicka, R.M. Cooke, and D.A. Ababei. Mining and visualising ordinal data with non-parametric continuous BBNs. *Computational Statistics and Data Analysis*, 54(3):668–687, 2010.

[31] A. Sklar. Fonctions de réparation à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris*, 8:229–231, 1959.

[32] H. Joe. *Multivariate Models and Dependence Concepts*. Chapman & Hall, London, 1997.

[33] Régis Lebrun and Anne Dutfoy. An innovating analysis of the nataf transformation from the copula viewpoint. *Probabilistic Engineering Mechanics*, 24(3):312 – 320, 2009.

[34] F.L.M. Diermanse and C.P.M. Geerse. Correlation models in flood risk analysis. *Reliability Engineering and System Safety*, 105:64 – 72, 2012. ISSN 0951-8320.

[35] Régis Lebrun and Anne Dutfoy. Do rosenblatt and nataf isoprobabilistic transformations really differ? *Probabilistic Engineering Mechanics*, 24(4):577–584, 2009.

[36] D. Kurowicka and R.M. Cooke. *Uncertainty Analysis with High Dimensional Dependence Modelling*. Wiley, 2006.

[37] O. Morales Nápoles, D. Worm, and B. Dillingh. Framework for probabilistic scale transition in physico-chemical modeling of asphalt. *TNO report project number 034.24789*, 2011.

[38] Rudolf O Weber and Peter Talkner. Some remarks on spatial correlation function models. *Monthly Weather Review*, 121(9):2611–2617, 1993.

[39] J. Whittaker. *Graphical Models in applied multivariate statistics*. John Wiley and Sons, Chichester, 1990.

[40] Adrian FM Smith and Alan E Gelfand. Bayesian statistics without tears: a sampling–resampling perspective. *The American Statistician*, 46(2):84–88, 1992.

[41] Walter R Gilks. *Markov chain monte carlo*. Wiley Online Library, 2005.

[42] Zdenek P. Bažant and Er-Ping Chen. Scaling of structural failure. *Applied Mechanics Reviews*, 50(10):593–627, 1997. 10.1115/1.3101672.

[43] Erik H Vanmarcke. On the distribution of the first-passage time for normal stationary random processes. *Journal of applied mechanics*, 42(1), 1975.

[44] Erik Vanmarcke. *Random fields: analysis and synthesis*. The Massachusetts Institute of Technology, 1983.

[45] David S Bowles, S Sanjay, Loren R Anderson Chauhan, and Ryan C Grove. Baseline risk assessment for herbert hoover dike. In *ANCOLD Conference on Dams*, 2012.

[46] United States Bureau of Reclamation. Dam safety risk analysis: Best practices and risk methodology training manual. Technical report, U.S. Bureau of Reclamation, 2011.

[47] Ruben Jongejan, Bob Maaskant, Wouter ter Horst, Fred Havinga, Niels Roode, and Harry Stefess. The vnk2-project: a fully probabilistic risk analysis for all major levee systems in the netherlands. *IAHS-AISH publication*, pages 75–85, 2013.

[48] Willem Kanning. *The weakest link: spatial variability in the piping failure mechanism of dikes*. TU Delft, Delft University of Technology, 2012.

[49] W Kanning, T Schweckendiek, T Vrouwenvelder, and E Calle. The length-effect in the reliability of flood defenses (in preparation). in prep.

[50] T Schweckendiek, ACWM Vrouwenvelder, and EOF Calle. Updating piping reliability with field performance observations. *Structural Safety*, 47:13–23, 2014.

[51] Robert E Melchers. Safety and risk in structural engineering. *Progress in Structural Engineering and Materials*, 4(2):193–202, 2002. ISSN 1528-2716.

[52] Robert H Sues and Mark A Cesare. System reliability and sensitivity factors via the mppss method. *Probabilistic Engineering Mechanics*, 20(2):148–157, 2005.

[53] A Naess, BJ Leira, and O Batsevych. System reliability analysis by enhanced monte carlo simulation. *Structural safety*, 31(5):349–355, 2009.

[54] A Naess, BJ Leira, and O Batsevych. Reliability analysis of large structural systems. *Probabilistic Engineering Mechanics*, 28:164–168, 2012.

[55] Won-Hee Kang and Junho Song. Evaluation of multivariate normal integrals for general systems by sequential compounding. *Structural Safety*, 32(1):35–41, 2010.

[56] Junho Chun, Junho Song, and Glaucio H Paulino. Parameter sensitivity of system reliability using sequential compounding method. *Structural safety*, 55:26–36, 2015.

[57] Michael Hohenbichler and Rudiger Rackwitz. Non-normal dependent vectors in structural safety. *Journal of the Engineering Mechanics Division*, 107(6):1227–1238, 1981.

[58] M Hohenbichler and R Rackwitz. First-order concepts in system reliability. *Structural safety*, 1(3):177–188, 1983.

[59] S Gollwitzer and R Rackwitz. Equivalent components in first-order system reliability. *Reliability Engineering*, 5(2):99–115, 1983.

[60] M Hohenbichler, S Gollwitzer, W Kruse, and R Rackwitz. New light on first-and second-order reliability methods. *Structural safety*, 4(4):267–284, 1987.

[61] RB Jongejan and B Maaskant. Applications of vnk2, a fully probabilistic risk analysis for all major levee systems in the netherlands. *Comprehensive flood risk management. Taylor & Francis Group, London*, 2013.

[62] JM Kind. Economically efficient flood protection standards for the netherlands. *Journal of Flood Risk Management*, 7(2):103–117, 2014.

[63] Armen Der Kiureghian et al. First-and second-order reliability methods. *Engineering design reliability handbook*, pages 14–1, 2005.

[64] M Van Koningsveld, GJ De Boer, F Baart, T Damsma, C Den Heijer, P Van Geer, and B De Sonnevile. Openearth-inter-company management of: data, models, tools & knowledge. In *Proceedings WODCON XIX Conference: Dredging Makes the World a Better Place, 9-14 September 2010, Beijing, China*, 2010.

[65] Ove Ditlevsen, Robert E Melchers, and H Gluver. General multi-dimensional probability integration by directional simulation. *Computers & Structures*, 36(2):355–368, 1990.

[66] RE Melchers. Structural system reliability assessment using directional simulation. *Structural Safety*, 16(1):23–37, 1994.

[67] Robert E. Melchers. *Structural reliability analysis and prediction*. John Wiley & Son Ltd, 1999.

[68] Carel Eijgenraam, Jarl Kind, Carlijn Bak, Ruud Brekelmans, Dick den Hertog, Matthijs Duits, Kees Roos, Pieter Vermeer, and Wim Kuijken. Economically efficient standards to protect the netherlands against flooding. *Interfaces*, 44(1):7–21, 2014.

[69] A Jeuken, J Kind, and J Gauderis. Cost-benefit analysis of flood protection strategies for the rhine-meuse delta. *Comprehensive flood risk management: research for policy and practice*, 228, 2012.

[70] Alan Genz and Frank Bretz. Numerical computation of multivariate t-probabilities with application to power calculation of multiple contrasts. *Journal of Statistical Computation and Simulation*, 63(4):103–117, 1999.

[71] Alan Genz and Frank Bretz. Comparison of methods for the computation of multivariate t probabilities. *Journal of Computational and Graphical Statistics*, 11(4): 950–971, 2002.

[72] TAW. Technical report on sand boils (piping). Technical report, Technical Advisory Committee on Flood Defences (Dutch), 1999.

[73] JB Sellmeijer and MA Koenders. A mathematical model for piping. *Applied mathematical modelling*, 15(11-12):646–651, 1991.

[74] Hans Sellmeijer, Juliana López de la Cruz, Vera M van Beek, and Han Knoeff. Fine-tuning of the backward erosion piping model through small-scale, medium-scale and ijkdijk experiments. *European Journal of Environmental and Civil Engineering*, 15(8):1139–1154, 2011.

[75] K Roscoe, S Caires, F Diermanse, and J Groeneweg. Extreme offshore wave statistics in the north sea. *WIT Transactions on Ecology and the Environment*, 133, 2010.

[76] ISO2394 ISO. General principles on reliability for structures. *International Standard Organization*, 2014.

[77] Randal Douc and Olivier Cappé. Comparison of resampling schemes for particle filtering. In *Image and Signal Processing and Analysis, 2005. ISPA 2005. Proceedings of the 4th International Symposium on*, pages 64–69. IEEE, 2005.

[78] Deltares. D-geo stability, slope stability software for soft soil engineering, 2014. URL https://www.deltares.nl/app/uploads/2014/11/D-Geo-Stability-web.pdf.

[79] Andrew Schofield and Peter Wroth. *Critical state soil mechanics*. McGraw-Hill London, 1968.

[80] IJkdijk Foundation. Ijkdijk macro-stability experiment, 2008. URL http://www.floodcontrolijkdijk.nl/downloads/ijkdijk_macro_stability_english.pdf.

[81] Stichting Deltawerken Online. The flood of 1953. URL http://www.deltawerken.com/The-flood-of-1953/89.html.

[82] JK Vrijling. Probabilistic design of water defense systems in the netherlands. *Reliability engineering & system safety*, 74(3):337–344, 2001.

[83] F Den Heijer and FLM Diermanse. Towards risk-based assessment of flood defences in the netherlands: An operational framework. *Electrical Measuring Instruments and Measurements*, page 243, 2012.

[84] Peter van den Berg. Binnen drie uur stond er bijna twee meter water in de wijk. watersnood bracht tuindorpers samen. *de Volkskrant*, January 14, 1995.

[85] G Apostolaksi. Probability and risk assessment: the subjectivistic viewpoint and some suggestions. *Nuclear Safety*, 19(3):305–315, 1978.

[86] Stanley Kaplan and B John Garrick. On the quantitative definition of risk. *Risk analysis*, 1(1):11–27, 1981.

[87] Karl N Fleming, Ali Mosleh, and R Kenneth Deremer. A systematic procedure for the incorporation of common cause events into risk and reliability models. *Nuclear Engineering and Design*, 93(2-3):245–273, 1986.

[88] GE Apostolakis. Uncertainty in probabilistic safety assessment. *Nuclear Engineering and Design*, 115(1):173–179, 1989.

[89] Jan van de Graaff. Probabilistic design of dunes; an example from the netherlands. *Coastal Engineering*, 9(5):479–500, 1986.

[90] J. van de Graaf. Probabilistic methods for dune design; Background of the TAW Guideline (in Dutch). Technical report, Delft University of Technology, Civil Engineering Department, Coastal Engineering Group, 1984.

[91] TAW. Guideline for the assessment of the safety of dunes as sea defense. (in Dutch). Technical report, Technische Adviescommissie voor de Waterkeringen, the Hague, Netherlands, 1984.

[92] R.J. Cirkel. Guideline for the design of river levees: Part 1: Upstream rivers (in Dutch). Technical report, TAW: Technische Adviescommissie voor de Waterkeringen, the Hague, Netherlands, 1984.

[93] R.J. Adringa. Guideline for the design of river levees: Part 2: Downstream rivers (in Dutch). Technical report, TAW: Technische Adviescommissie voor de Waterkeringen, the Hague, Netherlands, 1989.

[94] J Van Dixhoorn. Eastern scheldt storm surge barrier, proceedings of the delta barrier symposium, rotterdam 13-15 october 1982. 1982.

[95] CUR/TAW. Probabilistic design of flood defences. Technical report, Technische Adviescommissie voor de Waterkeringen and Center for civil engineering research and codes, 1990.

[96] Herman Van Der Most and Mark Wehrung. Dealing with uncertainty in flood risk assessment of dike rings in the netherlands. *Natural Hazards*, 36(1-2):191–206, 2005.

[97] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

[98] Erik Vanmarcke. *Random fields: analysis and synthesis*. World Scientific, 2010.

A

THE ISSUE OF CORRELATED COMPONENTS IN SYSTEM RELIABILITY

System reliability is notoriously complex to compute, due to the interdependence of system components. In levee systems, the interdependency arises from shared soil and load variables (for dependence between failure modes), and their spatial correlation (for dependence between segments).

In system reliability, components can be connected in different ways. In series, where if any component fail, the whole system fails, in parallel, where all components must fail for the system to fail, or a hybrid combination of these. In general, levee systems are series systems, because if any segment fails due to any failure mode, the system fails. Each component has two states, failure or survival.

The failure probability for a series systems, $P_{f,SS}$, is calculated as the probability of the union over all the component failure states. Conversely, the system reliability is the union over the component survival states, or $P_{s,SS} = 1 - P_{f,SS}$. If we let the state of failure for component i be denoted E_i, the failure probability for the system is expressed as shown in Equation A.1.

$$P_{f,SS} = P\left(\bigcup_{i=1}^{n} E_i\right) \tag{A.1}$$

From combinatorics, the inclusion-exclusion principle dictates how to combine sets. For the simple case of $n = 2$ components, we would get $P_{f,SS} = P(E_1) + P(E_2) - P(E_1 \cap E_2)$. For $n = 3$ components, $P_{f,SS} = P(E_1) + P(E_2) + P(E_3) - P(E_1 \cap E_2) - P(E_2 \cap E_3) - P(E_1 \cap E_3) + P(E_1 \cap E_2 \cap E_3)$. Figure A.1 shows the Venn diagram that illustrates an example of the inclusion-exclusion principle for three sets. For more components than 3, the series system failure probability can be written generally according to Eq. A.2.

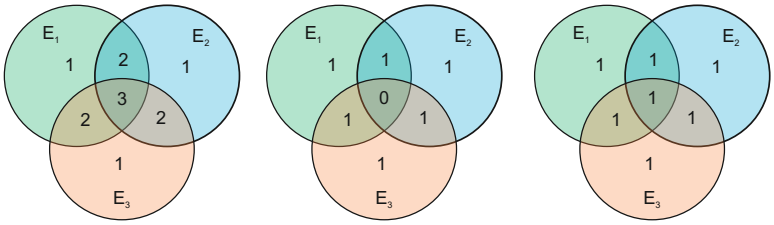

$E_1 + E_2 + E_3$

$E_1 + E_2 + E_3$
$-E_1 \cap E_2 - E_1 \cap E_3 - E_2 \cap E_3$

$E_1 + E_2 + E_3$
$-E_1 \cap E_2 - E_1 \cap E_3 - E_2 \cap E_3$
$+E_1 \cap E_2 \cap E_3$

Figure A.1: Example of the inclusion-exclusion principle for the union of sets, for three sets E_1, E_2, and E_3. The union is the combined area of these three sets. Adding the three sets together (left-hand picture) results in certain areas being double, or even triple counted (inclusion step). The second picture shows the subtraction of the areas $E_1 \cap E_2$, $E_1 \cap E_3$, and $E_2 \cap E_3$, which results in each area being counted once, except the area $E_1 \cap E_2 \cap E_3$, which is now excluded from the total (exclusion step). The last step (right-hand picture) is the inclusion of the missing area $E_1 \cap E_2 \cap E_3$.

$$P\left(\bigcup_{i=1}^{n} E_i\right) = \sum_{i=1}^{n} P(E_i) - \sum_{1\le i<j\le n} P\left(E_i \cap E_j\right) + \sum_{1\le i<j<k\le n} P\left(E_i \cap E_j \cap E_k\right) - \ldots + (-1)^{n-1} P(E_1 \cap \ldots \cap E_n) \tag{A.2}$$

Only in the trivial cases that all component states are fully correlated or entirely independent can Eq. A.1 be solved easily. If all component states are fully correlated, the series system failure probability is simply equal to the $\max(P(E_i))$. If they are independent, De Morgan's law can be used. For a two-component system this states that $E_1 \cup E_2 = \neg(\neg E_1 \cap \neg E_2)$, where $\neg$ is the negation logic operator (i.e. 'not'). So since E_i is the state of failure of component i, $\neg E_i$ signifies survival of component i. Then $P(E_1 \cup E_2) = P(\neg(\neg E_1 \cap \neg E_2)) = 1 - P(\neg E_1 \cap \neg E_2) = 1 - (1 - P(E_1))(1 - P(E_2))$, where the last step depends on the components being independent. For n independent components, De Morgan's law can be generalized according to Eq. A.3. In reality, the components of a levee system will never be fully correlated or independent, but rather will be partially correlated.

$$P_{f,SS;ind} = 1 - \prod_{i=1}^{n} 1 - P(E_i) \tag{A.3}$$

B

BASIC CONCEPTS IN LEVEE RELIABILITY

This dissertation is intended for readers with some experience in reliability analysis. However, this section provides a brief background to some of the concepts and notation dealt with in the dissertation.

In levee system reliability analysis, we are interested in the probability that a levee will fail. For the basic concepts, let us only consider a single component rather than a system. We are interested in the combinations of soil and load variables that will lead to a state of failure of the component. The probability of those combinations tells us the probability of failure. The state of failure is captured by a limit state function, denoted in this dissertation as Z. The limit state is defined such that it is positive when the levee survives, and negative when the levee fails. Consider a vector of random variables, $\mathbf{X} = [X_1, X_2, ..., X_n]$, containing both load and strength variables. The failure of the system is represented by the n-fold integral:

$$P_f = \int_{Z(\mathbf{X})<0} f_{\mathbf{X}}(X)\, dX \tag{B.1}$$

where $f_{\mathbf{X}}(X)$ is the multivariate density function of X. Eq. B.1 can rarely be evaluated analytically. A number of reliability methods are available to estimate the integral. Some commonly used ones are Monte Carlo (MC) sampling, MC Importance sampling, first-order reliability method (FORM), second-order reliability method (SORM), adaptive response surfaces, subset simulation, and numerical integration. Each method has its own advantages and disadvantages. Of these methods, I will briefly discuss MC and FORM, as they are most referred to in the dissertation.

CRUDE MONTE CARLO

Figure B.1 demonstrates a Monte Carlo sampling approach to evaluating the integral in Eq. B.1. To describe the method, let us consider a simple example where $\mathbf{X}$ consists of two random variables, $\mathbf{X} = [X_1, X_2]$, and where we have a limit state function

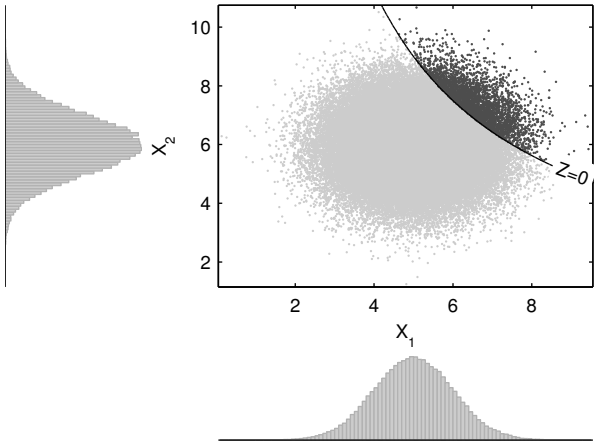

Figure B.1: Example of failure space. Samples of variables X_1 and X_2 are shown as dots, with their marginal distributions indicated by histograms. The limit state is shown as a solid black line. Combinations of X_1 and X_2 that lead to failure by a darker gray color.

$Z(\mathbf{X}) = C - X_1 \cdot X_2$, where X_1 is normally distributed with mean 5, and standard deviation 1 (denoted $X_1 \sim \mathcal{N}(5,1)$), $X_2 \sim \mathcal{N}(6,1)$, and $C = 45$. The figure illustrates a cloud of $[X_1, X_2]$ samples, as well as the limit state where $Z(\mathbf{X}) = 0$ (shown as a solid black line). The samples which lead to failure are indicated in the figure with darker points. The MC estimate of failure is the ratio of samples lying in the failure space (where $Z(\mathbf{X}) < 0$) to those in the survival space (where $Z(\mathbf{X}) > 0$).

$$P_f = \frac{1}{N} \sum_{i=1}^{N} \mathrm{I}(Z(\mathbf{X}_i) < 0) \tag{B.2}$$

FIRST-ORDER RELIABILITY METHOD

First-order reliability method (FORM) is an efficient algorithm to calculate the failure probability, but it is less exact than MC (which always converges to the correct value, given enough samples). FORM works in a standard normal space, where all variables are represented by independent standard normally distributed variables. In that space, the method searches for the point closest to the origin which results in $Z < 0$. This point is known as the design point, and has the highest density in the failure space. The method is efficient because it linearizes the limit state when it searches for the design point. At each iteration in the search procedure it numerically estimates the partial derivatives of the limit state $\partial Z / \partial X_i$, which specify the hyperplane (linearized limit state function). The method can then analytically determine the location where $Z = 0$ closest to the origin. The method then repeats the iteration at the estimated design point from the previous iteration, until convergence is reached. Figure B.2 shows the FORM approximation for the same example we used above to illustrate MC. It is clear that the FORM failure probability estimate will be an overestimate in this case, because it approximates a larger failure space. The MC estimate with 10^7 samples estimates $P_f = 0.038$, while

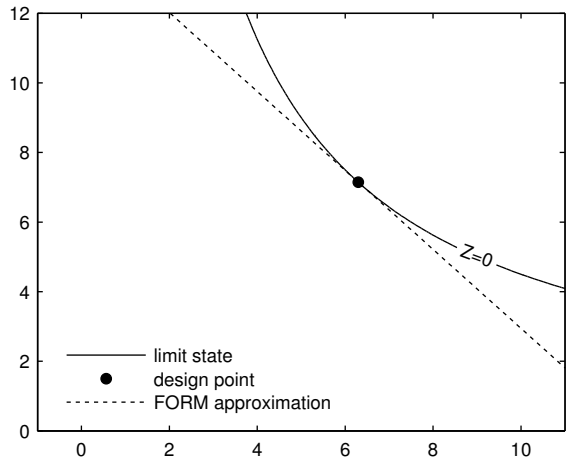

Figure B.2: FORM approximation of the limit state function for $Z = 45 - X_1 \cdot X_2$

FORM estimates $P_f = 0.042$.

The advantage of MC is that it is more accurate than FORM, if sufficient samples are taken. Unlike FORM, the nonlinearity of the limit state function has no impact on the MC estimate. The disadvantage of MC is that it requires a large number of samples to reach a sufficient accuracy. From [66], the number of samples, N, required to reach an acceptable relative error in the failure probability estimate, ε, can be calculated according to Eq. B.3, where k is the quantile of the standard normal distribution, such that the $\Phi(k) - \Phi(-k)$ returns the confidence that the actual error will be less than ε. For example, for 95% confidence, $k = 1.96$.

$$N = \frac{k^2}{\varepsilon^2}\left(\frac{1 - P_f}{P_f}\right) \tag{B.3}$$

For example, consider a failure probability of $P_f = 0.01$, an acceptable relative error of $\varepsilon = 0.05$, and a confidence of 95% ($k = 1.96$). The required number of samples would be about 150,000. For a more accurate estimate with $\varepsilon = 0.01$, N would need to be almost 4 million. For lower failure probabilities, the number of samples needed also increases substantially. If the limit state function is evaluated analytically, using a formula, a large value of N is not prohibitive. However, if evaluating the limit state requires calls to a geotechnical failure model, the computational burden quickly becomes insurmountable. For example, with $P_f = 0.001$, $\varepsilon = 0.05$, and $k = 1.96$, N is a little more than 1.5 million. If each sample called a failure model that required only 1 second to compute, the calculation time would be 17 days. Most failure models will require more than 1 second, making MC an unfeasible reliability method.

FORM usually only requires a handful of iterations before it converges to the design point. In the example presented in Figure B.2, only 20 iterations were needed. For each iteration, the partial derivatives must be computed, which requires a few evaluations of the limit state function. For a 1-sided evaluation of the partial derivatives, the number

of evaluations will be $(n+1) \cdot n_{iter}$. For a 2-sided evaluation, the number doubles. In the example in B.2, that translates to 60 or 120 calls to the limit state function. If the same hypothetical failure model is used, that would translate to a maximum of two minutes to compute the failure probability, compared with 17 days. That is the strong advantage of the FORM method.

A useful by-product of FORM is a set of influence coefficients of the random variables, which convey the relative influence each variable has on the limit state, and on the failure probability. This is useful information, because if we know a variable is very influential, we can divert more resources to better estimate it (e.g. more soil borings). The influence coefficients are also used heavily in the spatial and temporal upscaling of failure probabilities in the levee system reliability model used in Dutch national flood defense assessments. There are methods to derive influence coefficients when using MC, but these must be implemented separately. One of these, the center-of-gravity method, is explained in Appendix F.

C

RELIABILITY MODELING IN THE NETHERLANDS: A BRIEF HISTORY

The history of flood protection in the Netherlands is rich and complex, dating back centuries. In this appendix, the scope is limited to the history that is relevant to the development of flood defense reliability modeling. Providing this history should put into context the problems addressed in the dissertation.

In 1953, the approach to flood protection and reliability modeling changed rather drastically. In the night between January 31st and February 1st, a storm over the North Sea caused record-breaking water levels along the Dutch coast. The results in Zeeland and South Holland were catastrophic. In total, 89 levees were breached, and 1836 people lost their lives ([81]). The inundated area is indicated in Figure C.1.

Figure C.1: Flood extent, 1953 flood in the Netherlands

The tragedy shocked the nation, and resulted in the government setting up the Delta Commission on February 18, 1953, which was a group of 14 experts - 12 civil engineers,

Figure C.2: Levee breach at Den Bommel in South Holland, as a result of the 1953 storm. Image source: https://beeldbank.rws.nl, Rijkswaterstaat

an agronomist, and an economist - whose role was to advise the government on flood protection to ensure such a flood would never happen again. The Delta Commission changed the way levees were designed and built. Prior to the 1953 storm, levees were built with a height of the highest previously recorded water level plus some additional freeboard ([82], [83]). The Delta Commission considered this a reactionary approach, and moved towards a proactive approach in which the levees needed to be designed to withstand an extreme water level, with a specific exceedance probability. The Commission implicitly made this approach risk-based by assigning the required exceedance probability according to the economic activity of a protected area and the density of the inhabitants. The members of the committee envisioned a risk-based approach to designing and managing the levees, but a lack of knowledge and limited computing power prohibited it from being realized at that time.

In 1965 the Technical Advisory Commission for Flood Defenses (acronym: TAW in Dutch) was set up in part as a consequence of the 1960 flood in Tuindorp Oostzaan, an area just outside of Amsterdam. A leaking water pipe led to the collapse of a levee, and the subsequent filling of the polder. According to the newspaper The Volkskrant, the entire polder filled 2 m deep within 3 hours ([84]). The task of the TAW was to provide requested or unrequested advice with regard to all technical aspects relating to flood defenses directly to the minister.

In the 1970s and 1980s there was a surge in probabilistic method development. Internationally, this research was predominantly focused on nuclear risk assessment ([85], [86], [87], [88]). In the Netherlands, the TAW was hard at work developing probabilistic methods for levee ([7]) and dune ([89]) design. The TAW was critical of the methods then used in practice. They noted that failure can occur due to a number of failure mechanisms, and not only the high water levels which were considered at the time. And they emphasized that the probability that a levee section can withstand an extreme water level does not equate to the probability that the protected area will flood. In 1979, TAW set up a special group whose focus was probabilistic methods, with the goal of implementing a new philosophy of levee design, that could calculate the flooding probability. They envisioned their output as a guidance document - a sort of recipe book - for designing flood defenses. By the mid-80s, elements of the probabilistic methods were brought together in guidance documents developed for practical use in designing dunes ([90], [91], and river levees ([92], [93]). The Eastern Scheldt barrier was being designed and built using probabilistic methods ([94]); it was the first hydraulic structure to be probabilistically designed. In 1985, the TAW presented a comprehensive state-of-the art of probabilistic modeling for flood defense design at that time (five years later much of the material was published in English, see [95]). They noted that it was far from complete, and that many knowledge gaps needed to be filled in. Specifically, the methods for combining correlated components, and for estimating the length effect, needed improvement.

In the 1990s, research into flooding probabilities of levee systems accelerated, and computing power also improved. By 2000, the TAW published a report 'From exceedance probability to flooding probability'. Their report provided a new method, the result of a research project begun in 1992, that realized the vision of the Delta Commission: it could calculate the flooding probability of a protected area. In their research, they used

the then-newly-developed reliability model PC-Ring. Methods for combining correlated components, and for estimating the length effect, had been developed and incorporated in the model. They applied their method and software to four 'dike rings' (connected flood defenses (mostly levees) and higher ground which enclose a protected low-lying area). They proposed that flooding probabilities of the remaining 50 dike rings in the Netherlands be carried out in the following years.

Starting in 2001, the project National Flood Risk Analysis in the Netherlands (Dutch acronym VNK, English acronym FLORIS) was set up with the task to compute the flooding probabilities and flood risk of each of the dike rings in the Netherlands ([96]). While the TAW indicated this would take two years, in reality it took about 13 years, and was only completed in 2014. The results were used to inform new risk-based safety standards ([62]), which passed in the Dutch House of Representatives ('tweede kamer') in July 2016. The Dutch senate ('eerste kamer') approved the new standards on November 1, 2016, and they went into effect on December 20, 2016.

For the assessment of flood defenses according to the new risk-based standards, The Ministry of Transport, Public Works and Water Management commissioned the development of new software, Hydra-Ring, that was based on PC-Ring. The new software needed to be more transparent, more flexible, well-documented, and contain additional functionality so that it could be used in an operational setting for the assessment of flood defenses under the new standards, which are based on flooding probabilities. The software development began in 2009, and is largely completed, but still ongoing in 2016. The algorithms and model structure of Hydra-Ring can be seen as a culmination of over 60 years of research.

D

D-SEPARATION IN BAYESIAN NETWORKS

This appendix provides an informal, intuitive explanation of D-separation. For more detailed information, the reader is referred to [97]. D-separation (where the "D" stands for dependence) is a way of determining whether two variables are dependent or independent, both marginally and conditional on observed variable(s), by studying the arrows in the network. Specifically, two variables are independent (D-separated) if there are no active paths between them. In this appendix, I give the rules for determining what constitutes an active path. When a variable is observed, it can change an inactive path to an active one, and vice versa. Thus, observing variables changes the independence/dependence among variables.

Let us start by considering three variables, A, B, and C, where A and B are connected through C. There are four possible networks in this case, enumerated below. In each network, I specify whether the path is active (influence flows between A and B) or inactive, both for the case that C is unobserved or observed. When C is observed, that is indicated by drawing a circle around the variable C. Following these 'rules' is an intuitive explanation about why these paths are active or inactive.

1. A is an indirect cause of B:

   $A \rightarrow C \rightarrow B$: Active path

   $A \rightarrow \textcircled{C} \rightarrow B$: Inactive path

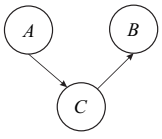

2. B is an indirect cause of A:

$A \leftarrow C \leftarrow B$: Active path

$A \leftarrow \textcircled{C} \leftarrow B$: Inactive path

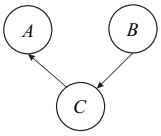

3. C is a common cause of A and B:

$A \leftarrow C \rightarrow B$: Active path

$A \leftarrow \textcircled{C} \rightarrow B$: Inactive path

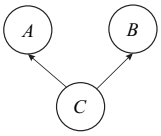

4. C is a common effect of A and B:

$A \rightarrow C \leftarrow B$: Inactive path

$A \rightarrow \textcircled{C} \leftarrow B$: Active path

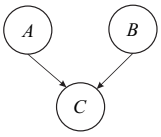

WHEN C IS NOT OBSERVED

In the first and second case, the flow of influence (dependence) from A to B and B to A is rather straightforward. For example, in case 1, A could be rainfall, C could be flood level (e.g. height of flooded water at a particular spot in a city), and B could be damage to property. The more likely the rainfall is to be heavy, the more likely the flood level is to be high, and the more likely the damage is to be high, and vice versa.

In the third case, C is a common cause of A and B. Influence in this structure can also pass from A to B. Consider the example that C is the flu, A is a fever, and B is a sore throat. If it is very likely that a person has a fever, it is more likely they have a flu, which in turn makes it more likely they have a sore throat.

In the fourth case, C is a common effect of A and B. In this case, there is no influence passed from A to B. Consider another rainfall example, where A is rainfall, B is poor drainage, and C is flood level. In this case, a high probability of heavy rainfall increases the probability a high flood level, but does not tell us anything about the drainage.

So to summarize the four cases when C is not observed, information can pass from A to B and from B to A in the first three cases. This is called an active path. In these cases, C is referred to as a *non-collider*, and A and B are considered D-connected. In the fourth

case, which is sometimes referred to as a V-structure (where two arrowheads come into a node), C blocks information from A to B, and is referred to as a *collider*; in this case, A and B are D-separated, and the path from A to B is an inactive path.

When C is observed

When C is observed, the flow of influence is actually just the opposite as when C is unobserved. This means that in the first three cases, A and B are independent (D-separated), given C, whereas in the fourth case A and B are dependent (D-connected) given C. Let's consider the examples looked at before. In case 1, the example was: A is rainfall, C is flood level, and B is property damage. If the flood level is observed, then the probability of the rainfall has no more influence on the property damage. It exercised its influence on the property damage by influencing the probability of flood level. Knowing C in this case breaks the flow of influence from A to B. In case 4, where A is rainfall, B is poor drainage, and C is flood levels, knowing C allows influence to pass from A to B. For example, prior to observing the flood level, a light rainfall told us nothing about the drainage in the city. However, once we know that the flood level was high, the information that the rainfall was light means that the drainage must be bad. Thus, when C is known, the probability of A influences the probability of B, and vice versa.

In summary, when non-colliders are unobserved, influence passes through them. When they are observed, they block influence. Conversely, colliders block influence when they are unobserved, and allow influence through when they are observed. Though not shown in this 3-node example, whenever a descendant of a collider is observed, it also activates the path through the collider.

Example

A more complex graph is shown in Figure D.1, to provide an example.

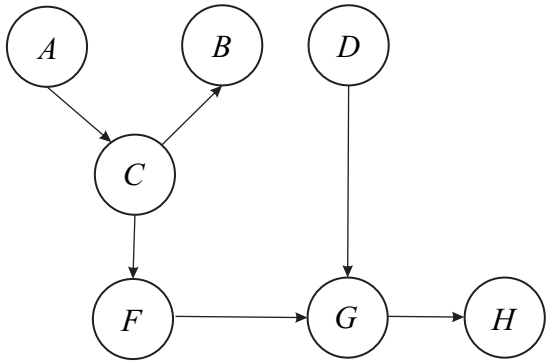

Figure D.1: Example BN for explaining D-separation (dependence and independence among variables which can be read from the graph).

The example contains seven nodes: A, B, C, D, F, G, and H. All of the nodes are non-colliders except for G. In the enumerated list below, I present the dependence statements about node A and all the other nodes in the network, for the case that (1) no nodes are observed, (2) node F is observed, and (3) node H is observed.

1. No nodes are observed:
   - A is d-connected to $\{B, C, F, G, H\}$
   - A is d-separated from D
2. Node F is observed:
   - A is d-connected to $\{B, C\}$, given F
   - A is d-separated from $\{D, G, H\}$, given F
3. Node H is observed:
   - A is d-connected to $\{B, C, DF, G\}$, given H
   - A is d-separated from $\{\ \}$, given H

When no nodes are observed, A is D-connected to every other node via an active path, except for D, which is separated from A by the collider G. When F is observed, it breaks the connection between A and the descendants of F, so that A is only D-connected to B and C. When H is observed, as a descendant of the collider G, it activates the path from A to D.

To summarize this appendix, the graphical structure of a Bayesian network allows an analyst to know which variables in the network are independent, both with and without the inclusion of evidence, using the concept of D-separation. Without the network representation, it would be challenging to determine all of the posterior dependences among variables, after including evidence.

E

DETAILS OF THE MODIFIED OUTCROSSING METHOD

The modified outcrossing (MO) method begins with the failure probability of a cross section, $P_{f,CS}$, or similarly, the reliability index $\beta = \Phi^{-1}\left(P_{f,CS}\right)$. If we assume that the crossings are a Poisson process, the survival probability of the segment $P_{s,seg}$ is described by Eq. E.1, where ν_+ is the upcrossing rate (the rate that the limit state crosses zero from positive to negative, given a reliability index β), L is the length of the segment, and b is the width of a cross section, hereafter assumed to be negligible ($b \approx 0$). Note that the upcrossing rate is equal to half the crossing rate ν.

$$P_{s,seg} = \left(1 - P_{f,CS}\right) \exp\left(-(L-b) \cdot \nu_+(\beta)\right) \tag{E.1}$$

To solve Eq. E.1, the crossing rate is needed, which can be derived analytically if the limit state function is approximated as a one-dimensional Gaussian random field: $Z = \beta - U$, where U is a standard normally distributed variable. This follows directly when first order reliability method (FORM) is used to calculate the cross sectional failure probability. In fact, FORM is preferred, because it also returns influence coefficients of the load and resistance variables, which are needed (see below).

Following the work of Vanmarcke ([44]), the crossing rate for a stationary, ergodic, Gaussian process can be expressed as a function of $\dot{U}$ (the derivative of U) according to Eq. E.2.

$$\nu\left(\beta\right) = \phi\left(\beta\right) \cdot \mathrm{E}\left(\left|\dot{U}\right|\right) \tag{E.2}$$

where ϕ is the standard normal density function. $\dot{U}$ depends on the spatial autocorrelation of Z, the expression for which is given in Eq. E.3 ([10],[48]), where d_x is the correlation length, and ρ_x is the residual correlation at large distances. These two parameters depend on the autocorrelation of the load and resistance variables (see Eqs. E.4 and E.5)([10]), where α_i is the influence coefficient, $\rho_{x,i}$ is the residual correlation, and $d_{x,i}$ is the correlation length of the i^{th} variable.

$$\rho_Z(x) = (1-\rho_x)\exp\left(-\frac{x^2}{d_x^2}\right) + \rho_x \tag{E.3}$$

$$\rho_x = \sum_{i=1}^{n} \alpha_i^2 \rho_{x,i} \tag{E.4}$$

$$\frac{1}{d_x^2} = \frac{1}{1-\rho_x}\sum_{i=1}^{n} \alpha_i^2 (1-\rho_{x,i})\frac{1}{d_{x,i}^2} \tag{E.5}$$

The problem with the autocorrelation in Eq. E.3 is that it means the limit state is not ergodic, which was a requirement to use the upcrossing rate in Eq. E.2. Ergodicity assumes that any sample of a process should have the same mean as the ensemble of all possible samples, and this is not the case when the residual correlation ρ_x is not equal to zero. To account for this, the MO method separates the ergodic and non-ergodic parts of the limit state function. It then computes the ergodic part of the segment failure probability $P_{f,seg}(w)$, which is conditional on a value of the non-ergodic variable $W = w$. Subsequently it uses the theorem of total probability to obtain the total segment failure probability.

Accounting for the non-ergodicity means that instead of $Z = \beta - U$, we get Eq. E.6, where U is the ergodic part, and W is the non-ergodic part, both of which are standard normally distributed variables.

$$Z = \beta - U\sqrt{(1-\rho_x)} - W\sqrt{\rho_x} \tag{E.6}$$

For the ergodic part of Z, the upcrossing rate is given in Eq. E.7, which is derived from Eq. E.2 taking into account that $\dot{U}$ is standard normally distributed. In Eq. E.7, $\sigma_{\dot{U}}$ is the standard deviation of the variable $\dot{U}$, an expression for which is available via ([98]); see Eq. E.8.

$$\nu_+(\beta) = \frac{\sigma_{\dot{U}}}{2\pi}\exp\left[-\frac{\beta^2}{2}\right] \tag{E.7}$$

$$\sigma_{\dot{U}}^2 = -\sigma_U^2 \cdot \rho_Z''(0) = -\rho_Z''(0) \tag{E.8}$$

The variable $\rho_Z''(0)$ is the second derivative of the autocorrelation function, evaluated for a lag of zero. Making use of the correlation function in Eq. E.3 (with ρ_x set to zero for the ergodic part of Z), we can calculate the expression for $\sigma_{\dot{U}}$ in Eq. E.8, and combine this with Eq. E.7 to derive an expression for the upcrossing rate wherein all the variables are known; see Eq. E.9.

$$\nu_+(\beta^*(w)) = \frac{1}{\sqrt{2}\pi d_x}\exp\left[-\frac{\beta^*(w)^2}{2}\right] \tag{E.9}$$

The reliability index $\beta^*(w)$ is the reliability index for the ergodic part of Z and is conditional on $W = w$. The expression for $\beta^*(w)$ is given in Eq. E.10; it is derived from Eq. E.6.

$$\beta^*(w) = \frac{\beta - w\sqrt{\rho_x}}{\sqrt{1-\rho_x}} \tag{E.10}$$

The conditional survival probability of the segment, $P_{s,seg}(w)$, is given by Eq. E.11.

$$P_{s,seg}(w) = \left(1 - P_{f,CS}\right) \exp\left[-(L) \cdot \nu_+\left(\beta^*(w)\right)\right] \tag{E.11}$$

Filling in the expression for $\nu_+(\beta^*(w))$ in Eq. E.11, the conditional failure probability of the levee segment can be computed as $P_{f,seg}(w) = 1 - P_{s,seg}(w)$; see Eq. E.12.

$$P_{f,seg}(w) = 1 - \left(1 - P_{f,CS}\right) \exp\left[-\frac{L}{\sqrt{2\pi}d_x} \exp\left[-\frac{\beta^*(w)^2}{2}\right]\right] \tag{E.12}$$

To calculate the total failure probability of the segment, the method uses the theorem of total probability; see Eq. E.13.

$$P_{f,seg} = \int_W P_{f,seg}(w)\, f(w)\, dw \tag{E.13}$$

F

RELIABILITY UPDATING WITH THE MODIFIED OUTCROSSING METHOD

The modified outcrossing (MO) method has not been used in conjunction with failure probability updating based on survival observations. To compare the posterior segment failure probabilities of the MO and BN methods I needed to make some implementation choices. This appendix describes those choices.

The first step is to update the *cross-sectional* failure probability based on a survival observation, and then apply the MO method to scale it up to the failure probability of the segment. The inference at the cross-sectional level is performed using MC rejection sampling, similar to the BN (see Chapter 2, Section 2.3.5), but for a single cross section. This is the same method that Schweckendiek et al. describe in [50] for updating at the cross-sectional level (see Section 2.5 of that paper). The inference results in an empirical joint posterior density over the resistance variables, $f_{\mathbf{R}}^{post}$ (in the numerical example in Chapter 3, this would be the joint density over R_1 and R_2). Eq. F.1 describes the posterior failure probability; it is the integration (over the failure space $Z < 0$) of the joint density of all the random variables. In the numerical example in Chapter 3, the resistance and the load are independent, so that the joint density is the product $f_{\mathbf{R}}^{post} \cdot f_S$. I evaluate the integral in Eq. F.1 with MC sampling. It is not possible to use FORM to calculate the posterior failure probability because the joint posterior distribution of the resistance variables has a dependence structure that is difficult/impossible to capture in a parametric way (see Section 2.3.5).

$$P_f^{post} = \int_{Z(\mathbf{r},s)<0} f_{\mathbf{R}}^{post}(\mathbf{r})\, f_S(s)\, d\mathbf{r}ds \tag{F.1}$$

The MO method requires influence coefficients of all the random variables, which it uses to estimate the parameters of the autocorrelation function of the limit state (see Eqs. E.4 and E.5 in Appendix E). MC simulation does not automatically return influence coefficients the way that FORM does, so I used a method known as center of gravity. Note

that this step in the implementation is important, and one in which errors can be introduced. For example, using a less robust method of estimating the influence coefficients can lead to large differences between MO and BN posterior segment failure probabilities that are not due to the MO method. I recommend the center of gravity method specifically because it is a robust (consistent) and accurate method.

The center of gravity method translates the posterior samples of the random variables (in the example in Section 3.3, this would be R_1, R_2, and S) to independent standard normal variables (U_{R1}, U_{R2}, and U_S). It then takes the mean over the samples which led to failure as the center of gravity. The method then searches the line between the center of gravity and the origin for the limit state (where $Z = 0$), and takes that point to be the design point, denoted $\left[u_{d_{R1}}, u_{d_{R2}}, u_{d_S}\right]$. The design point can be written in terms of the influence coefficients and the reliability index: $\left[u_{d_{R1}}, u_{d_{R2}}, u_{d_S}\right] = [\alpha_{R1}\beta, \alpha_{R2}\beta, \alpha_S\beta]$. The equality can be used to solve for the influence coefficients α_{R1}, α_{R2}, and α_S; specifically $\alpha_{R1} = u_{d_{R1}}/\beta$, $\alpha_{R2} = u_{d_{R2}}/\beta$, and $\alpha_S = u_{d_S}/\beta$.

Once the influence coefficients and reliability index for the cross section are derived, the MO method can be carried out as described in E.

G

CORRELATION WITHIN A CROSS SECTION

In the examples we consider in the dissertation, the resistance variables within a cross section are independent. In reality, these variables may be correlated. In this appendix, we provide details how the network will need to be modified to account for correlations among variables within a cross section.

We consider an example with two resistance variables, and two cross sections. We have the grain diameter d and the permeability of the soil k. These variables have been shown to be correlated (ref), and are relevant in the piping failure mechanism. The variables are also autocorrelated. The autocorrelation coefficients between k_1 and k_2 and between d_1 and d_2 are denoted ρ_k and ρ_d, respectively. The correlation between d and k is denoted ρ_{dk}, and this value is the same in each cross section. The network for the joint probability over d_1, k_1, d_2, and k_2 is presented in Figure G.1.

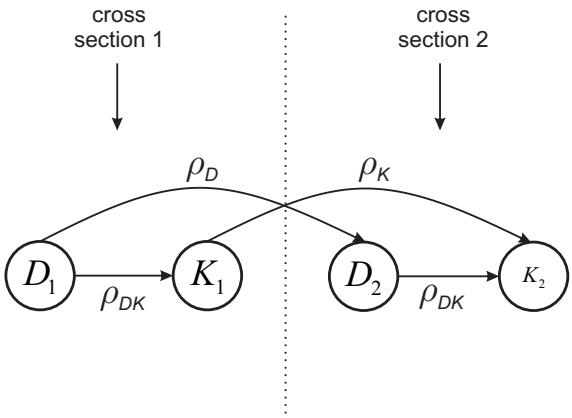

Figure G.1: Example of correlation within a cross section

The presentation in Figure G.1 helps visualize the network in terms of the physical cross sections. We present the same network in a different way, to better highlight the relationship between the four variables in Figure G.2. Shown next to the network in Figure G.2 is the correlation matrix for the four variables. If the resistance variables were assumed independent, the terms ρ_{dk} would be equal to zero, as would the correlations between d_1 and k_2 and between d_2 and k_1, and our entire correlation matrix would be known. However, when d and k are correlated within a cross section, it imposes correlation between d in one cross section with k in a neighboring cross section (see Figure G.2). While we are not explicitly interested in this correlation, we must include it to ensure that the matrix remains positive definite - which is a measure of consistency in the correlation structure. For example, for the two-cross section network in Figure G.2, consider the case where both ρ_{DK} and ρ_K are equal to 0.95. This means that a sample of D_1 that is much higher than average results in a K_1 sample that is much higher than average, and that influences the K_2 sample so that is also much higher than average. If we would assign a correlation of zero between D_1 and K_2, that would be impossible to satisfy, given the relationship between D_1 and K_1 and the relationship between K_1 and K_2.

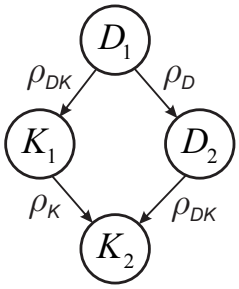

	D_1	K_1	D_2	K_2
D_1	1	ρ_{DK}	ρ_D	?
K_1	ρ_{DK}	1	?	ρ_K
D_2	ρ_D	?	1	ρ_{DK}
K_2	?	ρ_K	ρ_{DK}	1

Figure G.2: Network for the joint distribution over d_1, k_1, d_2, and k_2, with the associated correlation matrix. The correlation between d in one cross section with k in a neighboring cross section is generally unknown and represented by a question mark.

In the absence of dense soil boring data, the correlation between D in one cross section and K in a different cross section must be estimated based on the graph (i.e. not from data). We see in the example in Figure G.2 that there are two lines of influence between D_1 and K_2 and between K_1 and D_2. Let us consider D_1 and K_2. There is influence via K_1 (i.e. D_1 influences K_1 which influences K_2), and via D_2 (i.e. D_1 influences D_2 which influences K_2). If we consider each line of influence separately (that is, assuming the other line is absent), the correlation via K_1 would be $\rho_{DK} \cdot \rho_K$. This is explained by Equation G.1, which shows how we can sample from the 3-variable joint distribution (assuming a Gaussian copula). The variables are sampled in the standard Normal space. The first variable, U_{D1}, is sampled from a univariate standard Normal distribution. The second variable, U_{K1} is a combination of U_{D1} (the part of K_1 that is correlated with D_1),

and an independent standard Normal variable K_1^*. The third variable, U_{K2}, is a combination of U_{K1} and an independent standard Normal variable K_2^*. If we expand the expression for U_{K2} to include the expression for U_{K1}, we get the final line in Equation G.1. The influence that U_{D1} has on U_{K1} is given by the coefficient $\rho_K \cdot \rho_{DK}$ in the last line of Equation G.1.

$$\begin{aligned} U_{D1} &= U_{D1} \\ U_{K1} &= \rho_{DK} \cdot U_{D1} + \sqrt{1-\rho_{DK}^2} \cdot U_{K1}^* \\ U_{K2} &= \rho_K \cdot U_{K1} + \sqrt{1-\rho_K^2} \cdot U_{K2}^* \\ &= \rho_K \cdot \left[\rho_{DK} \cdot U_{D1} + \sqrt{1-\rho_{DK}^2} \cdot U_{K1}^*\right] + \sqrt{1-\rho_K^2} \cdot U_{K2}^* \end{aligned} \tag{G.1}$$

We can do the same analysis for the line of influence via D_2. In that case, the correlation between D_1 and K_2 would be $\rho_D \cdot \rho_{DK}$. The two lines of influence are both impacting the correlation between D_1 and K_2, but taking the maximum is sufficient to ensure that the correlation matrix is not inconsistent, and remains positive definite. Thus the cross-correlation term between D_1 and K_2, denoted ρ_{DK12}, is given in Equation G.2.

$$\rho_{DK12} = \max\left[\left(\rho_D \cdot \rho_{DK}\right), \left(\rho_K \cdot \rho_{DK}\right)\right] \tag{G.2}$$

EXAMPLE WITH THREE CROSS SECTIONS

We now expand the discussion to three cross sections. We revise the notation here a little bit for clarity; see Table G.1. The BN for three cross sections is shown in Figure G.3. Similar to the case with two cross sections, this presentation helps visualize the physical cross sections. Figure G.4 shows a more compact representation which helps us better visualize the connections between variables. Also shown in Figure G.4 is the correlation matrix for the Gaussian copula. Correlations between D_i and K_j for $i \neq j$ are unknown, and are represented in the matrix with question marks.

Table G.1: Explanation of correlation symbols in Figure G.3

Symbol	Correlation between
$\rho_{D\Delta1}$	D_i and D_j, $\lvert i-j \rvert = 1$
$\rho_{D\Delta2}$	D_i and D_j, $\lvert i-j \rvert = 2$
$\rho_{K\Delta1}$	K_i and K_j, $\lvert i-j \rvert = 1$
$\rho_{K\Delta2}$	K_i and K_j, $\lvert i-j \rvert = 2$
ρ_{DK}	D_i and K_i

We apply the same approach we used for a two cross-section BN to calculate the correlation between K in one cross section and D in a different cross section. We denote this correlation $\rho_{D_iK_j}$ for $i \neq j$. For example, let us consider the correlation $\rho_{D_1K_3}$. There are two routes of influence between D_1 and K_3 (see Figure G.4): (1) $D_1 \rightarrow D_3 \rightarrow K_3$ or (2) $D_1 \rightarrow K_1 \rightarrow K_3$. If we consider the first in the absence of the second, the influence of the first would be the product $\rho_{D\Delta2} \cdot \rho_{DK}$. Similarly, in the absence of the first, the influence of the second would be the product $\rho_{DK} \cdot \rho_{K\Delta2}$. We take the maximum of these two lines

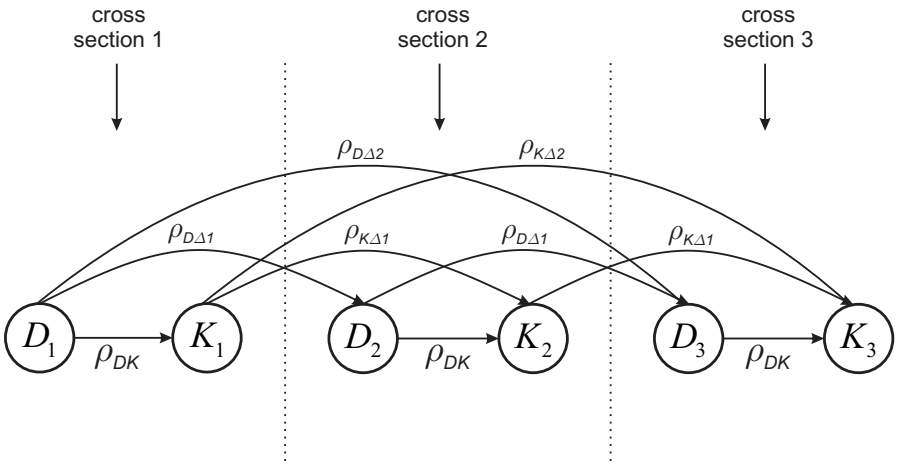

Figure G.3: Example of correlation within a cross section, for a three cross-section example

of influence to represent the correlation between D_1 and K_3 in the correlation matrix, to ensure the positive-definiteness of the matrix.

GENERALIZATION TO n CROSS SECTIONS

The line of reasoning in the previous sections can be generalized for an arbitrary number of cross sections. Equation G.3 shows the general formula to calculate the correlation $\rho_{D_i K_j}$ for $i \neq j$.

$$\rho_{D_i K_j} = \max\left[\left(\rho_{D\Delta|i-j|} \cdot \rho_{DK}\right), \left(\rho_{K\Delta|i-j|} \cdot \rho_{DK}\right)\right]; \; i \neq j \tag{G.3}$$

There are many degrees of freedom when trying to satisfy the positive-definiteness of a correlation matrix like the one in Figure G.4. The one we proposed is one of potentially many options. We tested it for up to 100 cross sections, and found that the matrix remains positive definite and the correlation matrix of the samples nearly perfectly agrees with the input matrix. For example, with 100 cross sections (so 200 variables, $K_1 ... K_{100}$ and $D_1 ... D_{100}$) and 1E6 samples, the largest difference between the sampled correlation matrix and the input correlation matrix was 3E-3.

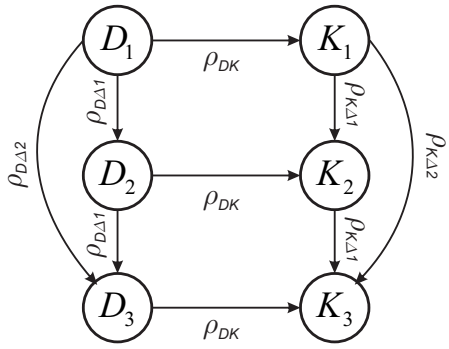

	D_1	D_2	D_3	K_1	K_2	K_3
D_1	1	$\rho_{D\Delta1}$	$\rho_{D\Delta2}$	ρ_{DK}	?	?
D_2	$\rho_{D\Delta1}$	1	$\rho_{D\Delta1}$	?	ρ_{DK}	?
D_3	$\rho_{D\Delta2}$	$\rho_{D\Delta1}$	1	?	?	ρ_{DK}
K_1	ρ_{DK}	?	?	1	$\rho_{K\Delta1}$	$\rho_{K\Delta2}$
K_2	?	ρ_{DK}	?	$\rho_{K\Delta1}$	1	$\rho_{K\Delta1}$
K_3	?	?	ρ_{DK}	$\rho_{K\Delta2}$	$\rho_{K\Delta1}$	1

Figure G.4: Network for the joint distribution over D_1, K_1, D_2, K_2, D_3, and K_3 with the associated correlation matrix. The correlation between D in one cross section with K in a neighboring cross section is generally unknown and represented by a question mark.

SAMENVATTING

Nederland wordt beschermd tegen overstromingen door een stelsel van waterkeringen. Het onderhoud en versterken van deze waterkeringen brengt hoge kosten met zich mee. Bij het maken van keuzes voor dijkversterkingen is het aan te bevelen om deze te baseren op gedegen risico-analyses. In dergelijke analyses moet onder andere vastgesteld worden hoe groot de kans is dat een waterkering gedurende de levensfase zal doorbreken. Het bepalen van deze faalkans, of overstromingskans, is complex en de uitkomsten zijn behept met onzekerheden. In de praktijk worden in overstromingskansberekeningen daarom vaak grove simplificaties toegepast.

In Nederland is de laatste decennia grondig onderzoek uitgevoerd met als uitkomst een geavanceerde rekenmethode voor het bepalen van overstromingskansen. Deze methode is geprogrammeerd in het software-pakket 'Hydra-Ring'. De rekenkern van Hydra-Ring bevat twee belangrijke algoritmen waarin (1) de faalkans van een dijkvak wordt bepaald en (2) de faalkansen van dijkvakken worden gecombineerd tot een totale systeemfaalkans. Het eerste algoritme, de 'Modified Outcrossing (MO) methode', brengt ruimtelijke correlaties van eigenschappen van de ondergrond, de dijkbekleding en het dijkprofiel in rekening. Het tweede algoritme, de 'Equivalent Planes (EP) methode', brengt correlaties tussen kenmerken van dijkvakken onderling in rekening. Beide algoritmen zijn efficiënt, maar een grondige beschrijving en verificatie ontbrak tot dusver. Verder zijn deze methoden nog niet toegepast in combinatie met informatie over 'bewezen sterkte', d.w.z. waarnemingen van hoogwaters die dijken hebben kunnen weerstaan. Deze dissertaties geeft invulling aan deze kennislacunes.

In de dissertatie wordt veelvuldig gebruik gemaakt van een techniek die bekend staat als 'Bayesian Networks' (BN). BNs zijn probabilistische modellen, waarmee correlaties tussen variabelen worden weergegeven in grafische structuren (netwerken). De BNs in deze dissertatie maken gebruik van Monte-Carlo (MC) sampling technieken om kansen te berekenen. Verder worden Gaussische copula-modellen gebruikt om de correlaties te kwantificeren. Op bepaalde onderdelen zijn de BNs met Gausissche copula-modellen exacter, d.w.z. minder 'benaderend' dan de algoritmen van Hydra-Ring en om die reden een excellent vehikel om de algoritmen van Hydra-Ring te valideren. In de regel vergen ze echter wel meer rekentijd.

De BN-modellen zijn toegepast in twee case-studies waarin onderzocht is wat de kans is op het optreden van het faalmechanisme 'piping'. In beide case-studies is gebruikt gemaakt van waarnemingen van hoogwaters waarbij de bewuste dijken niet hebben gefaald. Uit het onderzoek kwam naar voren dat de impact van deze waarnemingen op de geschatte faalkans van geval tot geval sterk kan verschillen. Op basis van een gevoeligheidsanalyse is vastgesteld onder welke condities dergelijke waarnemingen wel of juist geen significante impact hebben op de berekende faalkans.

Vervolgens is een BN ontwikkeld om de onzekerheid in de uitvoer van een geotechnische model te schatten. Deze onzekerheid heeft een dominante rol in de faalkans-

bepaling en het is daarom belangrijk om deze zo goed mogelijk af te schatten. In het onderzoek in deze dissertatie is aangetoond dat via het gebruik van de BN, waarnemingen van falen of niet falen van dijken, in combinatie met de voorspellingen op basis van het geotechnische model, een sterke verbetering kunnen opleveren van de schatting van de modelonzekerheid.

The BN-modellen die zijn ontwikkeld in het kader van deze dissertatie zijn een waardevolle toevoeging aan de vigerende methoden van betrouwbaarheidsanalyses van dijken. De BNs vergen significant meer rekentijd dan de benaderende algoritmen van Hydra-Ring, maar de rekentijden zijn over het algemeen acceptabel. Mede vanwege de verschillen in rekentijden komen BNs nog niet direct in aanmerking als vervanging van de algoritmen van Hydra-Ring. Echter, ze zijn zeer waardevol als verificatie voor de bewuste algoritmen. En indien waarnemingen van hoogwaters beschikbaar zijn waar dijken wel/niet hebben gefaald kunnen deze methoden gebruikt worden om de faalkansberekeningen aan te scherpen.

LIST OF PUBLICATIONS

PUBLICATIONS RELATED TO THE DISSERTATION

Roscoe, K., Hanea, A. and Vrouwenvelder, T., *Levee system reliability modeling: The length effect and Bayesian updating*, In review. Submitted to Structural Safety (2016).

Roscoe, K., Diermanse, F. and Vrouwenvelder, T., *System reliability with correlated components: Accuracy of the Equivalent Planes method*, Structural Safety **57**, pp.53-64 (2015).

Roscoe, K. and Hanea, A., *Bayesian networks in levee reliability*, ICASP12: 12th International Conference on Applications of Statistics and Probability in Civil Engineering, Vancouver, Canada (2015).

Diermanse, F., **Roscoe, K.**, IJmker, J., Mens, M. and Bouwer, L., *Hydra-Ring: a computational framework to combine failure probabilities*, EGU General Assembly Conference Abstracts, **15**, p. 11112 (2013).

PREVIOUS PUBLICATIONS

Roscoe, K., Weerts, A.H. and Schroevers, M., *Estimation of the uncertainty in water level forecasts at ungauged river locations using quantile regression*, International journal of river basin management **10(4)**, pp.383-394 (2012).

Groeneweg, J., Caires, S. and **Roscoe, K.**, *Temporal and spatial evolution of extreme events*, Coastal Engineering Proceedings **1(33)** (2012).

Roscoe, K. and Diermanse, F., *Effect of surge uncertainty on probabilistically computed dune erosion*, Coastal Engineering **58(11)**, pp.1023-1033 (2011).

Roscoe, K., Caires, S., Diermanse, F., and Groeneweg, J. *Extreme offshore wave statistics in the North Sea*, WIT Transactions on Ecology and the Environment **133** (2010).

Roscoe, K. and Diermanse, F., *Use of Projected Mean Discharge Values for Estimating Future Flood Risk*, EGU General Assembly Conference Abstracts, **11**, p. 8586 (2009).

Beckers, J.V.L., Sprokkereef, E. and **Roscoe, K.**, *Use of Bayesian model averaging to determine uncertainties in river discharge and water level forecasts*, Proc. 4th International Symposium on Flood Defence: Managing Flood Risk, Reliability and Vulnerability, Toronto, Ontario, Canada pp. 6-8 (2008).