OPTIMA 94

Mathematical Optimization Society Newsletter

MOS Chair's Column

April 15, 2014. Let us all welcome Volker Kaibel as the new Optima Editor-in-Chief. This is Volker's first issue, taking over after a great run by Katya Scheinberg. Joining Volker as Optima Co-Editors are Sam Burer and Jeff Linderoth. Thanks to Jeff, Sam, and Volker for keeping our society connected with Optima's lead articles, announcements, reports, and more.

The search for the Optima Editor-in-Chief completed a very busy year for the MOS Publications Committee. I would like to thank chair Nick Gould and his fellow committee members Darinka Dentcheva, Christoph Helmberg, Jie Sun, and Robert Weismantel for their long hours. Over the past year, the committee has carried out successful searches for the leadership positions of MPB, MPC, and the MOS/SIAM Book Series, together with Volker Kaibel's appointment. Allow me to make the announcements.

The editors of MOS publications are appointed to initial terms of four years, with a possible extension of an additional two years. The current Editor-in-Chief of MPA is Alex Shapiro, who began his term in the fall of 2012. For MPB, Jong-Shi Pang took over as Editor-in-Chief on January 1, 2014, following the long and successful run of Danny Ralph. For MPC, Dan Bienstock will take over as Editor-in-Chief on January 1, 2015. Last, but not least, Katya Scheinberg moved from Optima to take over as the Editor-in-Chief of the MOS/SIAM Book Series on January 1, 2014, following in the large footsteps of Tom Liebling.

Many thanks to all of our editors, past, present, and future!

Bill Cook, University of Waterloo
bico@uwaterloo.ca

Note from the Editors

Dear members of MOS:
It is our great pleasure to form the team of editors of our society's newsletter for the upcoming four years. Two of us (Sam and Volker) already enjoyed working as co-editors with editor Katya Scheinberg. Thanks to Katya for all the efforts she put into Optima and for being a very careful boss of the editorial team over the last three years.

From the feedback of many members of MOS we have the impression that Optima is for the most part appreciated in its current form. Thus, we do not see much need for major changes, but rather find ourselves faced with the challenge to keep the quality that the newsletter reached at the end of Katya's editorship (which to a great extent is due to the work of the entire earlier editorial team of Katya, Andrea Lodi and the late Alberto Caprara). We are more than happy that our new co-editor Jeff Linderoth is going to help us with this challenge!

It is probably not too overconfident to claim that the challenge is met with this current issue. We have an article by Ruud Brekelmans, Carel Eijgenraam, Dick den Hertog, and Kees Roos about their work on optimizing dike heights in the Netherlands, for which they received the 2013 INFORMS Franz Edelman Award for Achievement in Operations Research and the Management Sciences (see Page 7). The list of former winners of this prestiguous prize is quite impressive. The most famous felicitator probably was Nelson Mandela, whose letter of congratulations from 1996 still hangs in the INFORMS headquarters – in that year the prize went to the *South African National Defense Force*.

We are very grateful to Mike Trick for contributing the discussion column on the Edelman Prize and its relation to the field of Optimization.

We hope you will enjoy this issue, and we invite you to share with us any suggestions you might have, in particular with respect to covering specific topics in future issues!

Sam Burer, Co-Editor
Volker Kaibel, Editor
Jeff Linderoth, Co-Editor

Contents of Issue 94 / April 2014

Ruud Brekelmans, Carel Eijgenraam, Dick den Hertog and Kees Roos

A Mixed Integer Nonlinear Optimization Approach to Optimize Dike Heights in the Netherlands

1 Introduction

This article is based on the material in two earlier papers, Brekelmans et al. [2012] and Eijgenraam et al. [2014a] that were published in the journals *Operations Research* and *Interfaces*, respectively. It describes the optimization model that has been developed to optimize dike heights in the Netherlands. Moreover, it briefly describes the high impact of the results of this project on political decision making in the Netherlands. The project was awarded the INFORMS Franz Edelman 2013 Award. For more details on the validation of the model, the method used, and the political process and impact, we refer to these papers.

In the Netherlands, dike rings, consisting of dunes, dikes, and structures, protect a large part of the country against flooding.

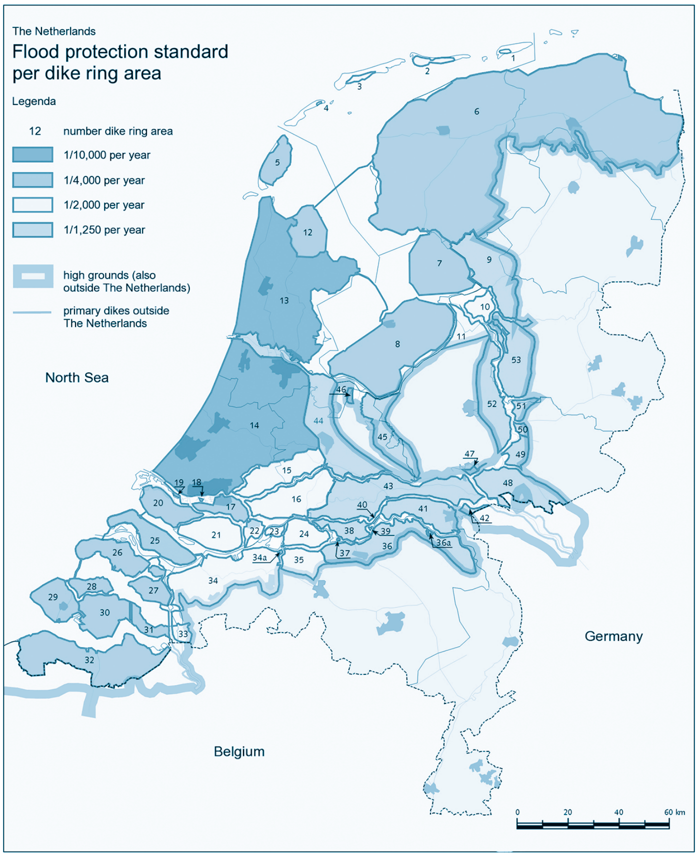

Figure 1. *53 main dike rings and the current legal protection standards*

After the catastrophic flood in 1953, a cost-benefit model was developed by D. van Danzig [1956] to determine optimal dike heights. The objective of the cost-benefit analysis (CBA) is to find an optimal balance between investment costs and the benefit of reducing flood damages, both as a result of heightening dikes. The question then becomes when and how much to invest in the dike ring. In Eijgenraam et al. [2014b] we improve and extend D. van Danzig's model. In that paper we show how to properly include economic growth in the cost-benefit model, and how to address the question when to invest in dikes. All these models consider dike rings that consist of a homogeneous dike. This means that all parts in the dike ring have the same characteristics with respect to investment costs, flood probabilities, water level rise, etc.

Many dike rings in the Netherlands, however, are nonhomogeneous, consisting of different segments that each have different characteristics. Differences occur, for instance, if along a dike ring in the delta area a river dominated regime changes into a sea dominated regime, or if a dike ring contains a large sluice complex. Currently, there are dike rings with up to ten segments in the Netherlands. In this nonhomogeneous case, it is not necessary and not desirable to enforce that all these segments are heightened simultaneously and by exactly the same amount. Hence, the decision problem for the nonhomogeneous case concerns when and how much to invest in each individual dike segment.

In the current article, we consider the extension of the homogeneous case in Eijgenraam et al. [2014b] to the nonhomogeneous case. The research has been carried out as part of a project initiated by the government. The project's main goal is to support decision-making with respect to setting new flood protection standards for the dike rings in the Netherlands. Figure 1 shows the main dike rings and the current legal protection standard for each dike ring. Efficient flood protection standards can be derived from the optimal investment strategy and the resulting flood probabilities. How this can be done is explained in Eijgenraam et al. [2014a]. Here we confine ourselves to a description of the first stage: finding the optimal investment strategy. In order to lay a firm base for the new standards, the 53 larger dike rings in the Netherlands need to be analyzed thoroughly. This requires that particular scenarios can be analyzed within a reasonable amount of time, where each scenario represents a certain instance of the model parameters such as economic growth, interest rate, water level rise, flood characteristics, investment costs and so on.

It is shown in Eijgenraam et al. [2014b] that the homogeneous case can be solved analytically. Unfortunately, we did not succeed in solving the nonhomogeneous case analytically. In this article we show how the nonhomogeneous dike height optimization problem can be modeled as a Mixed Integer Nonlinear Programming (MINLP) problem.

In addition to the MINLP formulation of the decision problem, we constructed an iterative optimization algorithm that speeds up the solution time considerably. The algorithm has been implemented in AIMMS, which has subsequently been integrated in user-friendly software to perform the dike ring analysis [Duits 2009a,b]. The final results have had a big impact in the political decision making process.

2 Nonhomogeneous Dike Height Optimization Problem

2.1 Problem Formulation

In this section we present our model for the nonhomogeneous dike height optimization problem. The model is an extension of the homogeneous problem, as introduced in Eijgenraam et al. [2014b]. The reader is referred to Eijgenraam et al. [2014b] for the foundation of the common model parts. A dike ring protects a certain area of land against flooding. The number of segments is denoted as L ($L \geq 1$). A dike ring is said to be nonhomogeneous if $L > 1$, and homogeneous otherwise. All segments can be heightened independently of each other. Moreover, each segment has its own properties with respect to investment costs and flood probabilities. To indicate the dependence of a model parameter on a particular dike segment, a subscript ℓ ($\ell = 1, \ldots, L$) will be added to this parameter. The set of all segments is denoted by $\mathcal{L}$.

The objective is to find an investment plan that minimizes the expected total costs. Only investments in a finite planning horizon $[0, T)$ are considered. An investment plan is represented by a tuple $(\mathbf{U}, \mathbf{t})$, with $\mathbf{U} \in \mathbb{R}_+^{L \times (K+1)}$ and $\mathbf{t} = (t_0, t_1, \ldots, t_K)^T$. The vector $\mathbf{t}$ represents the possible timings of dike segment heightenings, where $t_0 = 0 < t_1 < \cdots < t_K < T$. Hence, $K + 1$ is an upper bound on the number of segment heightenings in the planning horizon. For notational convenience, we denote $t_{K+1} = T$. The matrix $\mathbf{U}$ represents the segment heightenings, where the element $\mathbf{U}_{\ell k} = u_{\ell k}$ is the heightening (cm) of segment ℓ at time t_k ($\ell = 1, \ldots, L$, $k = 0, \ldots, K$). Of course, heightenings are assumed to be nonnegative. If $u_{\ell k} = 0$, then this means that segment ℓ is not heightened at time t_k. The ℓ-th row of $\mathbf{U}$, with the $K + 1$ heightenings of dike segment ℓ, is denoted by $\mathbf{u}^{(\ell)}$.

Throughout the remainder of this article we use the following notation for the cumulative segment heightening and the absolute segment height at time t ($t \geq 0$):

$$h_{\ell t} = \sum_{k: t_k \leq t} u_{\ell k}, \quad \text{and} \quad H_{\ell t} = H_{\ell 0}^- + h_{\ell t}.$$

where $H_{\ell 0}^-$ is the absolute height of segment ℓ immediately prior to a possible heightening at time $t = 0$. For notational convenience, we also use $h_{\ell k} = h_{\ell t_k}$ and $H_{\ell k} = H_{\ell t_k}$. Note that it follows from this definition that the segment height is a nondecreasing step function. Moreover, this implicitly means that heightenings are measured at the moment that the investment actions are completed. A lead time is not modeled.

The flood probability of segment ℓ at time t is given by

$$P_{\ell t} = P_{\ell 0}^{-}\,\mathbf{e}\left(\alpha_\ell(\eta_\ell t - h_{\ell t})\right), \tag{1}$$

with $\mathbf{e}(\cdot)$ the exponential function, $P_{\ell 0}^{-}$ (1/year) the initial flood probability, α_ℓ (1/cm) the parameter of the exponential distribution for extreme water levels and η_ℓ (cm/year) the structural increase of the water level. Both the hydraulic conditions and the quality of the dike segment are summarized by one indicator: height above the level that corresponds to the flood probability $P_{\ell 0}^{-}$. The weakest segment fully determines the flood probability of the entire dike ring. Hence, we define the flood probability of the entire dike ring at time t by $P_t = \max_{\ell\in\mathcal{L}} P_{\ell t}$.

A property that all segments have in common is that they protect the same area of land. Hence, if there is a flood, the damage does not depend on the segment in which a breach occurs. Furthermore, the potential damage costs increase in time with the economic growth rate γ. The damage costs do, however, also depend on the resulting height of the water level within a dike ring after a flood. In particular, along rivers the damage costs increase by the rise in the height of the lowest segment (in absolute height). Putting all this together yields the following damage costs, at time t, in the case of a nonhomogeneous dike ring:

$$V_t = V_0^{-}\,\mathbf{e}\big(\gamma t + \zeta(\min_{\ell\in\mathcal{L}} H_{\ell t} - \min_{\ell\in\mathcal{L}} H_{\ell 0}^{-})\big),$$

with V_0^{-} the initial damage costs and ζ (1/cm) the parameter that represents the increase in damage costs depending on the height of the lowest dike segment.

The expected damage costs at time t is given by the product of the flood probability and the damage costs:

$$S_t = P_t V_t = \max_{\ell\in\mathcal{L}} S_{\ell 0}^{-}\,\mathbf{e}\big(\beta_\ell t - \alpha_\ell h_{\ell t} + \zeta(\min_{\ell\in\mathcal{L}} H_{\ell t} - H_{\ell_0 0}^{-})\big), \tag{2}$$

where $S_{\ell 0}^{-} = P_{\ell 0}^{-} V_0^{-}$, $\beta_\ell = \alpha_\ell \eta_\ell + \gamma$ and $\ell_0 = \mathrm{argmin}_\ell H_{\ell 0}^{-}$. By using the fact that the segment heights remain unchanged in the interval $[t_k, t_{k+1})$, the total expected damage in this interval can be written as

$$\int_{t_k}^{t_{k+1}} S_t\,\mathbf{e}(-\delta t)\,dt = \mathbf{e}(-\zeta H_{\ell_0 0}) \int_{t_k}^{t_{k+1}} \mathbf{e}(-\delta t + \zeta \min_{\ell\in\mathcal{L}} H_{\ell t}) \max_{\ell\in\mathcal{L}} (S_{\ell 0}^{-}\,\mathbf{e}(\beta_\ell t - \alpha_\ell h_{\ell k}))\,dt, \tag{3}$$

where δ is the discount rate.

From an optimization point of view there are two problems with the integral in (3):

(i) The minimum absolute segment height $\min_\ell H_{\ell t}$ cannot be incorporated in an optimization model as a convex constraint.

(ii) Even though the segment heights do not change during the interval $[t_k, t_{k+1})$, the segment flood probabilities $P_{\ell t}$ as defined by (1) increase monotonically in time. Hence, the segment ℓ for which the maximum flood probability is obtained may change during the interval $[t_k, t_{k+1})$.

If we want to use (3) in a MINLP model, then we have to make some assumptions about these two issues. The minimum operator in (3) refers to the fact that the size of the damage depends on the segment that is lowest in absolute height. Since in practice it is usually clear which of the segments along rivers is the lowest in absolute height, it is assumed that this segment is known in advance. Let this dike segment be denoted by ℓ^*. It turns out that, for the dike rings in the Netherlands, this assumption is always satisfied.

An obvious approach to deal with the maximum operator in (3) is to interchange the integral and the maximum operator. Note that this yields a lower bound for (3), which introduces an error only if the segment for which the maximum is obtained changes within the interval $[t_k, t_{k+1})$. Clearly, the effect of the error will be more serious if the length of the interval is longer, and consequently this should be taken into account when defining the intervals. In the implementation of the MINLP model to be introduced in Section 2.2, we shall make sure that these intervals are small enough to guarantee a sufficiently accurate approximation.

Using the two assumptions from above, (3) can be approximated by

$$\mathcal{E}_k(\mathbf{U},\mathbf{t}) = \max_{\ell\in\mathcal{L}} \frac{S_{\ell 0}^{-}}{\beta_{1\ell}}\,\mathbf{e}\big(\zeta(H_{\ell^* t_k} - H_{\ell_0 0}^{-}) - \alpha_\ell h_{\ell k}\big) \left[\mathbf{e}(\beta_{1\ell} t_{k+1}) - \mathbf{e}(\beta_{1\ell} t_k)\right], \tag{4}$$

with $\beta_{1\ell} = \beta_\ell - \delta$. The total expected damage in the planning horizon $[0, T)$ is then approximated by

$$\mathcal{E}(\mathbf{U},\mathbf{t}) = \sum_{k=0}^{K} \mathcal{E}_k(\mathbf{U},\mathbf{t}).$$

Note that for a fixed investment plan, it is possible to evaluate the size of the approximation error, since we can accurately evaluate the minimum and maximum operators in (3). This evaluation can be used to obtain a true comparison between investment plans with different discretization schemes.

To take into account the period after the planning horizon, it is assumed that there are no changes to the expected damage after T, and hence no more investments are required. Thus, the discounted expected damage after the planning horizon is $S_T \int_T^\infty \mathbf{e}(-\delta t)\,dt$, which can be approximated analogously to (4), i.e.,

$$\mathcal{R}(\mathbf{U},\mathbf{t}) = \max_{\ell\in\mathcal{L}} \frac{S_{\ell 0}^{-}}{\delta}\,\mathbf{e}\big(\beta_{1\ell} T - \alpha_\ell h_{\ell K} + \zeta(H_{\ell^* t_K} - H_{\ell_0 0}^{-})\big). \tag{5}$$

The investment costs associated with the heightening of segment ℓ at time t_k depend, of course, on the actual amount of the heightening. The costs, however, are assumed to be independent of the heightening of other segments, regardless of the moments of these heightenings. We use the same investment cost function as introduced by Brekelmans et al. [2012], and refer to it as *exponential* investment costs. For any *positive* heightening $u_{\ell k}$, the exponential investment costs are given by

$$I_{\ell k}(\mathbf{u}^{(\ell)}) = (c_\ell + b_\ell u_{\ell k})\,\mathbf{e}\big(-\lambda_\ell \textstyle\sum_{i=0}^{k} u_{\ell i}\big), \qquad \mathbf{u}^{(\ell)} \in \mathbb{R}_+^{K+1}. \tag{6}$$

Hence, the investment costs depend on the amount of the heightening and the amount of the total heightening up to time t_k. Since there are no investment costs when there is no heightening, the investment cost function is discontinuous at zero, i.e.,

$$\mathcal{I}_{\ell k}(\mathbf{u}^{(\ell)}) = \begin{cases} I_{\ell k}(\mathbf{u}^{(\ell)}) & \text{if } u_{\ell k} > 0, \\ 0 & \text{if } u_{\ell k} = 0. \end{cases}$$

The total discounted investment costs in the planning horizon are then given by

$$\mathcal{I}(\mathbf{U},\mathbf{t}) = \sum_{\ell=1}^{L}\sum_{k=0}^{K} \mathcal{I}_{\ell k}(\mathbf{u}^{(\ell)})\,\mathbf{e}(-\delta t_k).$$

Since the objective is to minimize the sum of the investment costs and expected damage costs, the resulting optimization model can now be formulated as

$$\begin{aligned} &\min\ \mathcal{I}(\mathbf{U},\mathbf{t}) + \mathcal{E}(\mathbf{U},\mathbf{t}) + \mathcal{R}(\mathbf{U},\mathbf{t}) \\ &\text{s.t. } \mathbf{U} \in \mathbb{R}_+^{L\times(K+1)}, \quad t_0 = 0 < t_1 < \cdots < t_K < T. \end{aligned} \tag{7}$$

2.2 MINLP Model

This section discusses how the general dike height optimization problem (7) can be transformed into a mathematical optimization model that can be solved using optimization solvers. The problem as stated by (7) can be considered as a Non-Linear Programming (NLP) model since the decision variables **U** and **t** are continuous and the objective function's components are clearly nonlinear. From an optimization point of view, however, there are some issues that prevent us from actually solving the problem as stated by (7): the discontinuity of the investment cost functions at zero, and the approximation error of the expected damage in (4). The latter issue forces us to discretize the planning horizon, since continuous time variables could result in large intervals and consequently serious approximation errors. The discontinuity of the investment cost function can be resolved by discretization of the heightenings as well, or by adding binary decision variables that indicate whether a heightening is actually greater than zero or not. If both the moments and the amounts of the heightenings are discretized, then, theoretically, the problem can be solved using a dynamic programming approach. Unfortunately, the state space grows too large if multiple segments are considered, which implies that a dynamic programming approach is not applicable. Therefore, we consider a MINLP approach with discretization of the planning horizon only.

Next, the reformulation of problem (7) into a MINLP model is discussed. We assume that a *discretization scheme* $\mathbf{t} = (t_0, \dots, t_{K+1})$ with $t_0 = 0 < t_1 < \dots < t_K < t_{K+1} = T$ has been prefixed. The MINLP model then becomes:

$$\min \sum_{l,k} \mathbf{e}(-\delta t_k)(c_\ell y_{\ell k} + b_\ell u_{\ell k})\,\mathbf{e}(-\lambda_\ell \sum_{i=0}^{k} u_{\ell i}) + \sum_k E_k + R \tag{8a}$$

$$\text{s.t. } E_k \geq \frac{S_{\ell 0}^-}{\beta_{1\ell}} \mathbf{e}(\zeta d_k - \alpha_\ell h_{\ell k})\left[\mathbf{e}(\beta_{1\ell} t_{k+1}) - \mathbf{e}(\beta_{1\ell} t_k)\right], \tag{8b}$$

$$R \geq \frac{S_{\ell 0}^-}{\delta} \mathbf{e}(\beta_{1\ell} T - \alpha_\ell h_{\ell K} + \zeta(H_{\ell * K} - H_{\ell_0 0}^-)), \tag{8c}$$

$$h_{\ell k} = \sum_{i=0}^{k} u_{\ell i}, \tag{8d}$$

$$H_{\ell k} = H_{\ell 0}^- + h_{\ell k}, \tag{8e}$$

$$0 \leq u_{\ell k} \leq y_{\ell k} M, \quad y_{\ell k} \in \{0, 1\}, \tag{8f}$$

$$h_{\ell k}, H_{\ell k}, E_k, R \in \mathbb{R}, \tag{8g}$$

where, in (8b), we define $d_k = H_{\ell * k} - H_{\ell_0 0}^-$ simply to save horizontal space. Constraints (8b) and (8d)–(8g) are indexed over $\ell = 1, \dots, L$, $k = 0, \dots, K$, and constraint (8c) is indexed over $\ell = 1, \dots, L$. The objective function (8a) with sums over $\ell = 1, \dots, L$, $k = 0, \dots, K$ includes the exponential investment costs with the fixed cost component c_ℓ multiplied by $y_{\ell k}$. The binary variables $y_{\ell k}$ combined with (8f) are required to ensure that either $u_{\ell k} = 0$ and the investment costs in the objective function are zero, or $u_{\ell k} > 0$ and the investment costs are equal to $I_{\ell k}(\mathbf{u}^{(\ell)})$. In (8f), M denotes an upper bound of the highest possible dike heightening. The auxiliary variables E_k and R represent the expected damage costs in $[t_k, t_{k+1})$ and $[T, \infty)$ respectively. Constraints (8b) and (8c) are used to model the damage costs as convex constraints without using the maximum operator, as occurs in (4).

It is clear that the optimal solution to problem (8) is fully determined by the decision variables $u_{\ell k}$ ($\ell = 1, \dots, L$, $k = 0, \dots, K$). These decision variables could be considered the "pure" decision variables of problem (8), which, together with the discretization scheme **t**, represent the investment plan $(\mathbf{U}, \mathbf{t})$ that answers the fundamental questions of *when* and *how much* should be invested in dike heightening.

2.3 Implementation Issues

One of the project goals, set by the water-consultancy company Deltares, was that the model (8) could be solved for all major dike rings in a reasonable amount of time without the necessity to tune the algorithm's settings for specific dike rings. We were able to design a generic solution method that can solve any particular instance of the model without any fine-tuning. The model (8) has been implemented in AIMMS. Moreover, the software company HKV has integrated this model into the software package OptimaliseRing [Duits 2009a,b], used by the actual performers of the cost-benefit analysis. We used the AIMMS Outer Approximation (AOA) method that is implemented in AIMMS to solve the MINLP problems.

A heuristic algorithm is needed because the MINLP (8) cannot be solved exactly in reasonable time for dike rings with more than 6 segments. For example, we used a dedicated discretization scheme to reduce the number of variables, and we added (nearly) redundant constraints to reduce the search space. Moreover, since it is clear that model (8) requires the input of several parameters, which in practice are often uncertain, we also developed a regret approach to obtain a solution that is robust with respect to these uncertainties. For more details we refer to Brekelmans et al. [2012].

3 Numerical Results

As discussed in Section 2.3, the optimization algorithm has been implemented in AIMMS using the AOA solver. All numerical results in this section were obtained using AIMMS 3.8.5 with CPLEX 11.2 and CONOPT 3.14G on a PC with an Intel Core 2 CPU processor.

A database with data about the dike rings in the Netherlands was provided by Deltares. This database contains all relevant parameters for the nonhomogeneous dike height optimization problem.

Overview Dike Rings

A selection of the dike rings from Deltares' database were optimized by our optimization algorithm. For all experiments we used common values for the discount rate per year ($\delta = 0.0247$) and the economic growth rate per year ($\gamma = 0.019$). A summary of the results for the exponential investment costs is shown in Table 1. The first two columns give the dike ring number along with the number of segments in the dike ring. The third column gives the MINLP model's objective value of the algorithm's final iteration. The fourth column gives a true evaluation of this objective value that does not suffer from an approximation error in the expected damage. It can be seen that the MINLP's objective is indeed a lower bound and that the approximation error is very modest, which indicates that the approximation of the expected damage is suitable for our MINLP model.

The fifth column in Table 1 gives the solution time in minutes. There does not appear to be a clear relationship between the number of segments and the solution time. This is mainly due to the fact that the discretization scheme is created in such a way that the number of resulting decision variables does not depend on the number of segments. In other words, a dike ring with more segments has a rougher discretization scheme than a dike ring with less segments, as explained in Brekelmans et al. [2012].

For the same set of experiments, Tables 2–4 show the moments of the first three updates of the dike rings, which correspond to one or more segment heightenings taking place at the same point in time. In addition, the table shows the effect the heightenings have on the dike ring's flood probability, i.e., the flood probabilities just before and just after the update are listed. For the new safety standards in this example, there are five out of twelve dike rings that require immediate segment heightenings at $t = 0$. The results also

Table 1. *Results of the optimization algorithm for a selection of dike rings*

Dikering	*Segments*	MINLP objective (M€)	True objective (M€)	Solution time (min)
10	4	107.51	107.51	0.52
13	4	10.38	10.38	0.07
14	2	94.04	94.04	0.54
16	8	1044.45	1046.08	6.24
17	6	377.05	377.37	3.33
21	10	217.40	217.71	2.23
22	5	373.98	374.08	7.62
36	6	395.65	395.65	60.19
38	3	136.26	136.29	59.33
43	8	486.72	488.10	1.65
47	2	16.57	16.57	8.54
48	3	42.92	42.92	2.77

Table 2. *Moments (in years measured from the start of the planning horizon) of the first dike ring update and the flood probabilities just before* $(P^-(t))$ *and after* $(P^+(t))$ *the updates.*

	First heightening		
Dike ring	t	$P^-(t)$	$P^+(t)$
10	68	6.6×10^{-4}	6.7×10^{-5}
13	140	1.8×10^{-4}	1.6×10^{-5}
14	36	1.5×10^{-4}	2.3×10^{-5}
16	0	5.0×10^{-4}	2.8×10^{-4}
17	20	3.8×10^{-4}	9.1×10^{-5}
21	0	5.0×10^{-4}	2.5×10^{-4}
22	7	5.2×10^{-4}	4.5×10^{-5}
36	36	1.1×10^{-3}	1.7×10^{-4}
38	0	6.7×10^{-4}	2.7×10^{-4}
43	0	2.7×10^{-4}	2.7×10^{-4}
47	30	2.5×10^{-4}	1.2×10^{-5}
48	0	2.8×10^{-4}	1.2×10^{-5}

Table 3. *Moments (in years measured from the start of the planning horizon) of the second dike ring update and the flood probabilities just before* $(P^-(t))$ *and after* $(P^+(t))$ *the updates*

	Second heightening		
Dike ring	t	$P^-(t)$	$P^+(t)$
10	156	1.2×10^{-4}	1.4×10^{-5}
13	244	3.7×10^{-5}	2.5×10^{-6}
14	104	4.6×10^{-5}	6.5×10^{-6}
16	40	3.7×10^{-4}	7.7×10^{-5}
17	81	1.9×10^{-4}	1.3×10^{-5}
21	45	5.2×10^{-4}	5.3×10^{-5}
22	100	1.1×10^{-4}	8.5×10^{-6}
36	102	4.1×10^{-4}	6.3×10^{-5}
38	28	4.6×10^{-4}	1.9×10^{-5}
43	30	4.6×10^{-4}	3.9×10^{-5}
47	120	4.0×10^{-5}	1.2×10^{-5}
48	77	3.0×10^{-5}	2.9×10^{-6}

Table 4. *Moments (in years measured from the start of the planning horizon) of the third dike ring update and the flood probabilities just before* $(P^-(t))$ *and after* $(P^+(t))$ *the updates*

	Third heightening		
Dike ring	t	$P^-(t)$	$P^+(t)$
10	244	2.5×10^{-5}	2.9×10^{-6}
13	—	—	—
14	168	1.3×10^{-5}	1.8×10^{-6}
16	105	1.2×10^{-4}	2.5×10^{-5}
17	165	4.3×10^{-5}	2.9×10^{-6}
21	120	1.5×10^{-4}	1.4×10^{-5}
22	200	2.3×10^{-5}	1.2×10^{-6}
36	165	1.5×10^{-4}	2.4×10^{-5}
38	126	8.6×10^{-5}	3.2×10^{-6}
43	120	9.7×10^{-5}	7.3×10^{-6}
47	200	1.6×10^{-5}	5.8×10^{-7}
48	154	7.1×10^{-6}	6.6×10^{-7}

clearly indicate that the flood probabilities just prior to a heightening decrease over time. This is a result of the economic growth, which increases the damage costs if a flood occurs, and therefore it is beneficial to let the flood probabilities decrease over time.

Let us take a closer look at the resulting solution for a dike ring with 6 segments. Figures 2 and 3 give a graphical overview of the final solution obtained with the iterative algorithm. Figure 2 shows the cumulative heightenings of the six segments during the 300-year planning horizon. Figure 3 shows the resulting segment flood probabilities. It can be seen that the two segments *1* and *5* are not heightened together with the other segments at $t = 20$. Figure 3 also shows why it is not necessary to heighten these two segments: their flood probabilities are still very low compared to the other segments. Although in this particular example there is a moment at which not all segments are heightened simultaneously, the figure clearly demonstrates why simultaneity very frequently leads to very good, or even optimal, results. Recall that a dike ring's flood probability is determined by the maximum segment flood probability. Hence, if a single segment is not heightened simultaneously with the other segments, then it is likely that this segment's flood probability will become, or even remain, the dike ring's maximum flood probability. The benefit of heightening the other segments, in terms

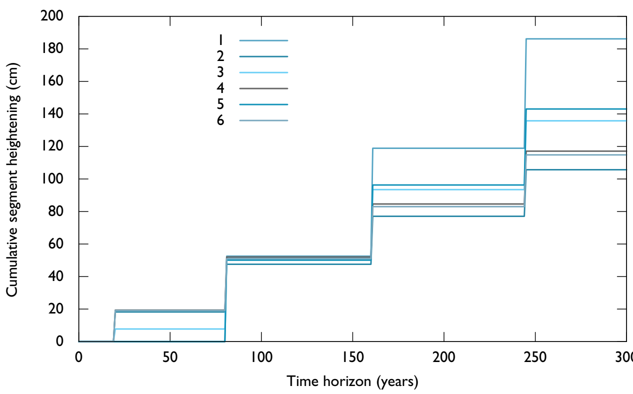

Figure 2. *Cumulative segment heightenings for a dike ring with 6 segments*

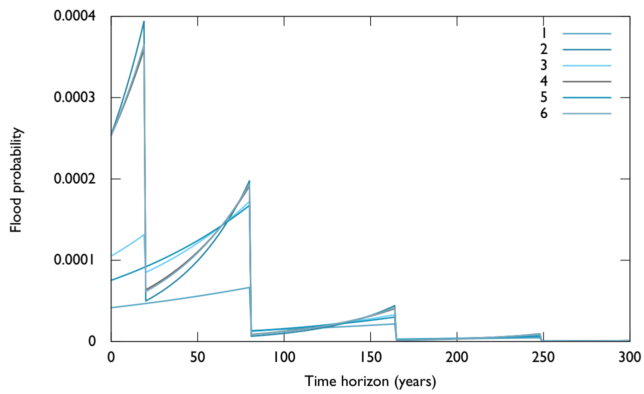

Figure 3. *Segment flood probabilities for the same dike ring as in Figure 2*

of decreasing the expected damage, is therefore usually smaller than the incurred investment costs.

4 Practical Impact

The ultimate goal of our project was to give recommendations for new flood protection standards in the Netherlands. In Eijgenraam et al. [2014a] it is described how to construct flood protection standards based on the optimal investment strategy resulting from our MINLP model. Based on the final results, published in Kind [2011, 2013], we concluded that increasing the legal protection standards of all dike-ring areas tenfold, as the Second Delta Committee recommended, is unnecessary. The current protection standards are (more than) appropriate, except for three regions: a part of the dike rings along the Rhine and Meuse Rivers (i.e., part of the areas that now have a standard of 1/2,000 or 1/1,250 per year), the southern part of dike ring 8 Flevoland (comprising the large, rapidly growing city of Almere), and some dike rings (e.g., 20) near Rotterdam.

The Water Advisory Committee, chaired by Crown Prince (currently the King) Willem-Alexander, discussed the final report of the CBA Kind [2011] and endorsed our results in a letter dated March 9, 2012. The House of Parliament discussed the report on December 5, 2011 and April 4, 2012. In a unanimous motion on April 17, 2012, parliament asked the government to present a concrete proposal for new legal standards in 2014, explicitly referring to the three regions named in Kind [2011, 2013] and under the condition that improvements are justified by a CBA. The state secretary of the Ministry of Infrastructure and the Environment (hereafter abbreviated as I&M) followed the report's results and recommendation: A tenfold increase in protection standards for all dike-ring areas is not needed and only the protection standards in the three regions named in the report need improvement.

The state secretary therefore instructed the Delta Commissioner to adapt, as necessary, the protection standards derived for these areas according to local situations, and to ensure that a minimal protection level is guaranteed everywhere in a dike ring area. On April 26, 2013, the Minister of I&M, Melanie Schultz van Haegen, confirmed these conclusions in a policy letter to the parliament.

The Delta Commissioner has announced that his proposal for new flood protection standards will closely follow the main conclusions of this project, which have already been recognized in discussions with the water boards and the provinces. In 2014, the cabinet will take a decision on these proposals. In 2015, the final decision on the improvement of these protection standards will be taken in parliament, such that new standards – after approval of the law in parliament – will be legally effective by 2017. Finally, in a letter dated November 27, 2012, the chairman of the renowned Second Delta Committee agreed with these conclusions, which clearly deviate from the committee's earlier advice. Compared to this earlier recommendation, this successful application of operations research yields both a highly significant increase in protection for these regions (in which two-thirds of the benefits of the proposed improvements accrue) and approximately 7.8 billion euro in cost savings.

Acknowledgement
We wish to thank all members of the project team "Optimal safety standards" for offering us their very useful knowledge and expertise. We are particularly grateful to Jarl Kind, Carlijn Bak (both Rijkswaterstaat and Deltares) and Matthijs Duits (HKV) for their extensive support. We would like to emphasize that this project was a multi-disciplinary project and that our model relies heavily on many research results carried out by the following organizations: CPB Netherlands Bureau for Economic Policy Analysis, Rijkswaterstaat, Deltares, HKV, KNMI, and Ministry of Infrastructure and the Environment. Without them, this project would not have been possible. We also wish to thank the editors of both *Operations Research* and *Interfaces* for getting permission to combine parts of the papers Brekelmans et al. [2012] and Eijgenraam et al. [2014a] in the current article.

Ruud Brekelmans, Dept. of Econometrics and Operations Research, Tilburg University, 5000LE Tilburg, The Netherlands. r.c.m.brekelmans@uvt.nl
Carel Eijgenraam, CPB Netherlands Bureau for Economic Policy Analysis, 2585JR The Hague, The Netherlands. c.j.j.eijgenraam@cpb.nl
Dick den Hertog, Dept. of Econometrics and Operations Research, Tilburg University, 5000LE Tilburg, The Netherlands. d.denhertog@uvt.nl
Kees Roos, Dept. of Information Systems and Algorithms, Delft University of Technology, 2628CD Delft, The Netherlands. c.roos@ewi.tudelft.nl

References

Brekelmans, R.C.M., D. den Hertog, C. Roos, and C.J.J. Eijgenraam. 2012. Safe dike heights at minimal costs: The nonhomogeneous case. *Operations Research*, 60(6):1342–1355.

Dantzig, D. van. 1956. Economic decision problems for flood prevention. *Econometrica*, 24(3):276–287.

Duits, M.T. 2009a. OptimaliseRing. Technische documentatie van een numeriek rekenmodel voor de economische optimalisatie van veiligheidsniveaus van dijkringen. Versie 2.0. HKV rapport PR1377, HKV LIJN IN WATER. In Dutch.

Duits, M.T. 2009b. OptimaliseRing. Testrapport. Versie 2.0. HKV rapport PR1377, HKV LIJN IN WATER. In Dutch.

Eijgenraam, C.J.J., J. Kind, C. Bak, R.C.M. Brekelmans, D. den Hertog, M. Duits, C. Roos, P. Vermeer, and W. Kuijken. 2014a. Economically efficient standards to protect the Netherlands against flooding. *Interfaces*, 44(1):1–15.

Eijgenraam, C.J.J., R.C.M. Brekelmans, D. den Hertog, and C. Roos. 2014b. Flood prevention by optimal dike heightening. Under review for *Management Science*.

Kind, J. 2011. Cost-benefit analysis water safety 21st century. Technical Report 1204144-006-ZWS-0012, Deltares, Delft, NL. In Dutch.

Kind, J. 2013. Economically efficient flood protection standards for the Netherlands. To appear in *Journal of Flood Risk Management*.

Michael Trick

The Franz Edelman Award and Optimization: A Wonderful Partnership

It is not immediately obvious that the Franz Edelman Award and the world of optimization have much in common, let alone a wonderful partnership as my title suggests. The Franz Edelman Award is an annual competition "to bring forward, recognize, and reward real-world applications of operations research, management science, and analytics". Optimization is a field of study that looks at techniques that find the best element within some mathematical structure. An Edelman presentation at an INFORMS conference is a *big deal*, with multiple presenters, fancy graphics, strict time limits, and an audience full of people in full business attire. An optimization talk at INFORMS? Not such a big deal. Presentations consist of slides full (or overfull) of Greek symbols, presented to an audience with attire upgraded from jeans and t-shirt, but not much. I have been known to remove a tie when presenting a particularly dense optimization talk, and I don't think I am the only one to do so. The papers that come out of Edelman presentations are discursive and wide-ranging; optimization papers are direct and to-the-point (and generally have an overabundance of Greek symbols).

Despite this "two-worlds" aspect, optimization and the Edelman Award are closely intertwined. Each greatly strengthens the other to the extent that each would be vastly different without the other. Without optimization, the Edelman Awards would be greatly lim-

Franz Edelman Award 2013, from left to right: Jarl Kind (Deltares), Dick den Hertog (University of Tilburg), Carel Eijgenraam (Netherlands Bureau for Economic Policy Analysis), Jaap Kwadijk (Deltares), Ruud Brekelmans (University of Tilburg), Kees Roos (Delft University of Technology)

ited in what they could recognize and reward. Without the Edelman Awards, the world of optimization would lose a key component that leads to the success of the field.

Before I address these two effects, let me review the Edelman Award competition.

I first encountered the Edelman Awards back in the late 1980s. I had just joined the business school at Carnegie Mellon University, and I was tasked to create a new course, entitled "Operations Research Applications". My instructions did not go much beyond that title: the school wanted a course that followed up on our core operations research courses, and that showed how operations research could be applied. After I learned more about MBA students, I decided that my initial approach of going through some of the easier chapters of Nemhauser and Wolsey's integer programming book was not likely to be successful. Carnegie Mellon MBAs are a technical bunch, but they are still MBAs. They want to know how a course will make them successful in business. So I built a course around the best source of application-oriented papers I could find: papers published in the journal *Interfaces*. Students read a variety of papers from that journal, wrote up summaries, and made presentations to the class. The class went great. Students were excited to learn that they could read and make sense of work in the professional literature, and *Interfaces* provided the right level of applicability so that students could see how this could affect their careers.

One thing I noticed, though, is that some students ended up presenting work that seemed much more developed than others from *Interfaces*. The problems seemed bigger, and the results seemed better validated. In short, the papers seemed more important. A bit more study led me to discover that all of these papers came from a January–February issue of the journal, and were from the Edelman Award competition. Non-Edelman papers were good for the class but Edelman papers were great!

Over the years, I added more Edelman papers to my "Operations Research Applications" course, until practically all the readings were Edelman papers. I began attending the Edelman presentations at the ORSA/TIMS and then INFORMS conferences. I watched videos of past Edelman competitions, and used the insights gleaned from those videos in the classroom. I began doing some consulting, trying to take what I had learned from the Edelmans and apply things in practice, combined with my own interests in optimization, of course. Over time, I was even asked to be a screener at the Edelman competition and later a judge in the final award selection process. I became an Edelman competition junkie.

The Edelman Award is really a competition. Every year, a call for applications goes out. This application is not extensive with just a four page limit. These applications (and there are generally 20–50 of them every year) are reviewed by a 35 member panel of screeners, who, in an enormous conference call, winnow them down to twelve or so semi-finalists. Every semi-finalist is assigned two or three verifiers who proceed to expand on and verify the claims that have been included in the application. This involves multiple calls and emails with the applicant, along with calls to clients, customers, and top executives to confirm the real-world effects of the application. If the applicant claims savings of a particular amount, the judges attempt to verify that number. Speculative or fanciful savings are quickly identified and dismissed, leaving only true, verifiable results.

Based on the verifiers reports, the full screening panel then reduces the group down to six or so finalists. Each finalist is assigned two or three coaches to aid them as they prepare their full paper and their hour-long presentation to be given at the spring INFORMS Analytics conference. Nine or so judges are chosen to form the final judging panel. To be successful, the finalists have to put a tremendous amount of work into their presentations. Those presentations often include things like videos of the CEO talking about the virtues of the project, customer commendations, and more, all trying to paint a full picture of the effect of the work. The winner is announced at a gala dinner, and all the finalists are inducted as Edelman Laureates.

This all seems a world away from, say, the COIN-OR Cup, an award with its own charm, but one that was handed out over beers around a pool table at a recent meeting.

How does optimization help the Edelman Award? Most obviously, many Edelman Award papers have, at their heart, an optimization problem to be solved. Whether it is finding the best way to route empty ships, or match up financial instruments, or allocate ads to a web site, there is some optimization that goes on. If you look over the last three years of finalists (18 in total) and check the keywords on the *Interfaces* articles, at least 12 have some form of optimization as a keyword. This includes integer programming (2 times), linear programming (3), network optimization (3), dynamic programming (2), stochastic programming, and such specialized terms as MINLP and branch-and-price. Not every Edelman finalist uses optimization (other methods such as simulation, statistical analysis, and system dynamics are also commonly used), but a good portion of the finalists have a signification optimization component.

Why is this? Why does optimization lead to Edelman success? I think there are a few reasons. First, the type of big, important, problems that are common in Edelman finalists lend themselves to the investment needed to use optimization. The difference between "optimal" and a heuristic solution within 2% of optimal takes on much more importance when the application involves hundreds of millions or billions of dollars. With millions on the line, it is worth the extra effort to get the best solution possible.

Second, large problems need formalism to provide boundaries on the scope of the work, and the Edelman award works within those boundaries to determine the true effect of the project.Optimization provides clear boundaries as to what is, and what is not, included in the project.If there are variables and constraints relative to part of the system, then that is in the project; otherwise it is not. This makes it easier to determine the impact of the project. Projects without these boundaries end up looking fuzzier, making it more difficult to compete for the Edelman Award.

Third, there is no doubt that the increase in the quality and quantity of data, combined with faster computers and better underlying optimization software greatly enhances the usefulness of optimization. When there are gigabytes (or more) of high-quality data, and a project can harness fast computers and excellent optimization software, projects have a way of being successful. The scope of optimization has grown tremendously in the last decade, making it more likely to show up in top projects of practical import.

So, if optimization is important to the Edelman Award, is the converse true? Is this a true partnership? Absolutely! There are many reasons why those in optimization should be interested in, and should support, the Edelman award.

The first, and perhaps most important, is the visibility the Edelman competition gets within an organization. A traditional part of an Edelman presentation is a video of a company CEO extolling the benefits of the project. While, in many cases, the CEO has already known about the project, this provides a great opportunity to solidify his or her understanding of the role of optimization in the success of the company. With improved understanding comes willingness to further support optimization within the firm, which leads to more investment in the field, which is good for optimization. As a side note, I find it a personal treat to watch CEOs speak of optimization with enthusiasm: they may not truly understand what they mean when they say "lagrangian based constrained optimization" but they can make a very convincing case for it.

The second area that the field can value is the identification of new directions for research. Recent Edelman finalists have highlighted the roles of robustness, uncertainty, vast amounts of data, nonlinearities, and other issues that point to exciting research directions for optimization researchers. By carefully looking at the Edelman papers, those in optimization can see not only what is used, but what could not be done due to limitations in the knowledge base of the field. This, in turn, leads to new, influential research. As an example, would the integer programming technique branch-and-price be so well developed were it not for the intense interest of airlines companies who use the technique for crew scheduling and other applications?

A final area I would like to highlight is the role the Edelman Award plays in attracting students to our field. My initial interactions with the Edelman Award came through my MBA class, and that class has continued for 25 years in attracting and inspiring students about operations research. Beyond the social good of having more people know about operations research, this has also solidified the role of operations research in our business school, making it more likely for our group to hire, support doctoral students, and otherwise be successful. Better courses mean more optimization at our school, and that is good for the field. Inspired students are also those likely to choose our field for doctoral study and contribute to the overall health of the field.

As a final example of the interaction between optimization and the Edelman Award, I need point no further than the paper "A Mixed Integer Nonlinear Optimization approach to optimize dike heights in the Netherlands" immediately adjacent to this article. This paper has all the signs of an optimization paper: it is full of Greek letters and theorems, and the rest. This paper provides the mathematical backing to the Edelman Award winning paper "Economically Efficient Standards to Protect the Netherlands Against Flooding" (*Interfaces*, 44, 7–21, 2014).

The *Interfaces* paper looks like an Edelman paper: it discusses politics, implementation, verifiable cost savings, and more. Together, they are much stronger than either alone.

While it may not be obvious that the world of optimization and the Edelman Awards work together, they most certainly do: successes in one mean success in the other. I look forward every January to the Edelman issue of *Interfaces* for there I see further signs of the robust health of the field of optimization.

Michael Trick, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA USA 15213. trick@cmu.edu

IMPRINT
Editor: Volker Kaibel, Institut für Mathematische Optimierung, Otto-von-Guericke Universität Magdeburg, Universitätsplatz 2, 39108 Magdeburg, Germany. kaibel@ovgu.de ■ **Co-Editors:** Samuel Burer, Department of Management Sciences, The University of Iowa, Iowa City, IA 52242-1994, USA. samuel-burer@uiowa.edu ■ Jeff Linderoth, Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison, WI 53706-1572, USA. linderoth@wisc.edu ■ **Founding Editor:** Donald W. Hearn ■ Published by the Mathematical Optimization Society ■ Design and typesetting by Christoph Eyrich, Berlin, Germany. optima@0x45.de ■ Printed by Oktoberdruck AG, Berlin, Germany.