Optimal Design and Life-Long Adaptation of Civil Infrastructure under Climate Change and Uncertain Demands

Ashmita Bhattacharya[1]; Gordon P. Warn[2]; Kostas G. Papakonstantinou[3]; Melissa M. Bilec[4]; Lauren McPhillips[5]; Chris E. Forest[6]; Rahaf Hasan[7]; Aditya Sharma[8]; and Digant Chavda[9]

[1]Dept. of Civil and Environmental Engineering, Penn State Univ. Email: azb6481@psu.edu
[2]Dept. of Civil and Environmental Engineering, Penn State Univ. Email: gpw1@psu.edu
[3]Dept. of Civil and Environmental Engineering, Penn State Univ. Email: kup31@psu.edu
[4]Dept. of Civil and Environmental Engineering, Univ. of Pittsburgh. Email: mbilec@pitt.edu
[5]Dept. of Civil and Environmental Engineering, Penn State Univ. Email: lxm500@psu.edu
[6]Dept. of Meteorology and Atmospheric Science, Penn State Univ. Email: cef13@psu.edu
[7]Dept. of Civil and Environmental Engineering, Univ. of Pittsburgh. Email: rah1930@pitt.edu
[8]Dept. of Civil and Environmental Engineering, Penn State Univ. Email: ams11335@psu.edu
[9]Dept. of Civil and Environmental Engineering, Penn State Univ. Email: dmc6948@psu.edu

Downloaded from ascelibrary.org by Gordon Warn on 11/30/23. Copyright ASCE. For personal use only; all rights reserved.

ABSTRACT

Various climate change effects pose increasing risks to the nation's infrastructure. Available methodologies address the risk-management problem primarily through cost-benefit analysis frameworks, which evaluate a comprehensive set of protection strategies against a wide range of simulated possible future scenarios. However, due to the substantial climate model uncertainties present over the future planning horizon, such strategies can often lead to less informed policies that might be optimal in an average sense, over the mean of anticipated future scenarios, but cannot offer adaptive solutions based on the actual climate effects evolving in time. To address these limitations, in this research, climate risk mitigation is instead formulated as a decision-making problem within a closed-loop stochastic control-based framework using Markov decision processes (MDP), taking real-time data into account, for evaluating the evolving conditions, and selecting the best possible, most informed life-cycle actions in time. Although broadly applicable, the merit of the framework will be illustrated through coastal risk mitigation against storm surge and sea-level rise in an idealized coastal city setting.

INTRODUCTION

The effects of intensified storms in conjunction with climate change induced sea level rise, heat waves, and other events, pose increasing risks to the nation's people and infrastructure. The increasing risks are, in part, due to the historical data, in which the nation's design codes and standards are based upon, no longer being representative of our changing climate. Risk-based, cost-benefit analysis frameworks can be used to evaluate a comprehensive set of protection strategies against an ensemble of representative potential future scenarios, eventually selecting the strategy with the minimum average life-cycle costs. However, due to the substantial climate model uncertainties present over the future planning horizon, such strategies can often lead to less informed design and adaptation decisions, that might be optimal in an average sense, over the mean of the simulated ensemble, but suboptimal based on the actual climate effects evolving in time. To overcome these challenges and limitations, there is an ongoing NSF LEAP-HI research project that aims to develop a framework to efficiently and simultaneously optimize

each phase of the life-cycle, while rigorously accounting for uncertainty in the life-cycle demands, system's state, and decision criteria. As part of this project, the overall framework and developments are being illustrated through application to the following climate risk mitigation cases: coastal flooding due to storm surge and sea-level rise, bridge scour due to riverine flooding, and building safety and energy consumption due to increasing demands associated with wind, temperature, and humidity.

This paper presents a snapshot and progress of the research project where the merit of the framework development is illustrated through a particular climate-related management application scenario involving coastal risk mitigation against storm surge and sea-level rise. This scenario is formulated as a sequential decision-making problem, in the form of a closed-loop stochastic control-based framework using Markov Decision Process (MDP), considering real-time data and all related evolving conditions and selecting the most informed, optimal life-cycle actions in time. The underlying environment model is developed based on historical sea-level observations and the most recent, sixth assessment report of the Intergovernmental Panel on Climate Change (IPCC) sea level rise projections, to quantify the overall rise long-term non-stationarity of sea level due to climate change. A suitable storm surge model is also devised. We solve the MDP models using point-based, asynchronous dynamic programming solvers that can provide globally optimal life-cycle solutions under uncertainty, directly suggesting distinct, optimal long-term actions at each time step. This developed control-based framework is demonstrated in an idealized coastal city setting, and its performance and capabilities are compared against existing techniques. Overall, the presented work aims to introduce a new paradigm of life-long optimal infrastructure adaptation based on observed data and realized conditions, considering and responding to the important uncertainties that characterize the problem.

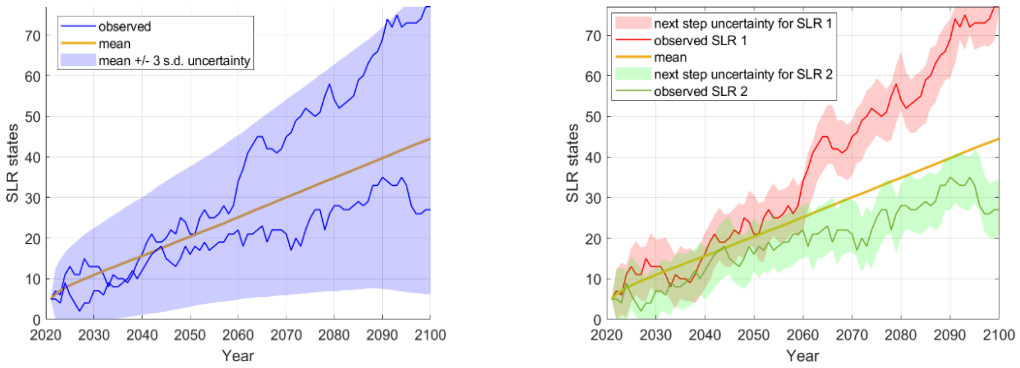

1(a) 1(b)

Figure 1. Illustration of two sea level rise (SLR) process decision-making options. Figure 1(a) Customary analysis with larger uncertainty growing over the future horizon, leading to same policy for the two trajectories shown. Figure 1(b) Framework presented herein, suggesting unique policies for the two trajectories and through periodic observations significantly reducing the band of uncertainty over the next time step from any current decision.

The current state-of-the-art methods in terms of cost-benefit analysis framework typically address the various uncertainties related to future flood events and damage estimations by simulating a large number of future events which are representative of the wide range of possible

Downloaded from ascelibrary.org by Gordon Warn on 11/30/23. Copyright ASCE. For personal use only; all rights reserved.

future scenarios (Ward et.al. 2017). In view of the substantial amount of uncertainty associated with all climate projection models (Arias et.al. 2021) over the future planning horizon, associated with civil infrastructure starting from the current conditions (Figure 1(a)), the cost-benefit analysis frameworks can lead to less-informed policies that overestimate or underestimate future climate needs. Since the strategies are evaluated on an ensemble of future climate scenarios, the optimal strategy chosen might be optimal over the mean of the ensemble, but the actual climate trajectory may deviate significantly from the mean, which makes the chosen strategy suboptimal with respect to the actual climate trajectory evolving in time. Due to the evolving nature of climate change effects, like sea level rise and other stressors contributing to flood hazards, the risk mitigation planning should ideally be adaptive.

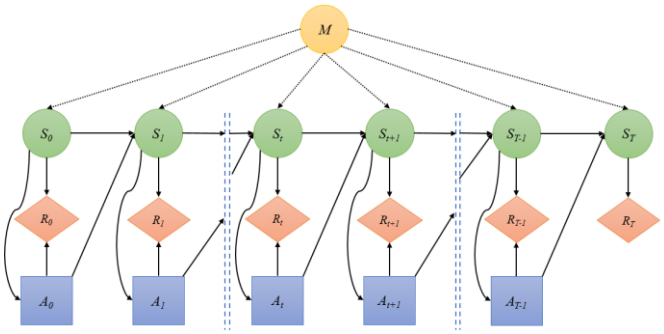

Figure 2. Sequential decision-making process as a Markov Decision Process. At each time step, the action taken is determined by the current state, the reward depends on the current state and the action taken, and the next state is determined by the Markovian transition dynamics of the underlying model given the current state, action pair.

To achieve adaptive risk-mitigation, this work presents the climate risk mitigation in the context of a MDP (Puterman 2014) stochastic control-based framework where decisions are made sequentially in time rather than primarily at the beginning of the planning horizon. MDPs consider real-time information and provide related optimal decisions, considering all possible future paths and alternatives, including non-stationary changes in climate and other sources. The acquired knowledge at each planning step greatly reduces the uncertainty present over the near future, before the next observation of the system becomes available, thus leading to more-informed life-cycle decisions with respect to the actual evolving climate conditions (Figure 1(b)). This work can also incorporate various new aspects into the relevant decision-making framework, such as carbon emission costs and the choice of green mitigation measures.

METHODOLOGY

The problem of determining optimal flood management decisions under uncertainty posed by climate change effects is herein addressed with methods falling under the broader category of stochastic control (Papakonstantinou and Shinozuka 2014), with emphasis on Markov Decision Processes (MDPs). A MDP is a controlled stochastic process represented by a 5-tuple $\langle S, A, P, R, \gamma \rangle$, where an agent observes the true state of the system $s \in S$, takes an action $a \in A$ from a finite set of possible actions A, receiving a reward $R(s,a)$, and as a result of taking this action, the system transitions to a new state s' according to a stochastic model with Markovian

Downloaded from ascelibrary.org by Gordon Warn on 11/30/23. Copyright ASCE. For personal use only; all rights reserved.

transition dynamics $P(s' \mid s, a)$, as illustrated in Figure 2. The long-term objective is expressed as an expected sum of discounted rewards generated from a given initial state and following a chosen policy, expressed as:

$$V_\pi(s) \coloneqq E_\pi \left[\sum_{k=0}^{\infty} \gamma^k R_{t+k} \mid S_t = s \right] \tag{1}$$

where this expectation in Eq. 1 provides the so-called value function. The state-dependent sequence of actions defines an agent's policy $\pi : S \to A$. The optimal value function V^* under the optimal policy π^* from a given system state *s* and for infinite horizon can be shown to follow the recursive Bellman equation (Puterman 2014):

$$V^*(s) = \max_{a \in A} \left\{ R(s,a) + \gamma \sum_{s \in S} p(s' \mid s, a) V^*(s') \right\} \tag{2}$$

This non-linear equation can be solved using dynamic programming algorithms such as value iteration or policy iteration. In this work, an asynchronous value iteration (VI) algorithm referred to as Focused Real Time Dynamic Programming (FRTDP) (Smith and Simmons 2006) is employed to solve the resulting MDP. The algorithm keeps track of both upper (V_U) and lower bounds (V_L) of the value function, such that $V_L \le V^* \le V_U$, which are used in the convergence criterion of the algorithm, and also help to guide the prioritized search over the state space.

APPLICATION: COASTAL CITY SETTING

To illustrate the adaptive climate risk mitigation framework in the context of flood management, an idealized coastal city is considered, inspired by Manhattan in the NYC area, as illustrated in Figure 3. The combined flood management and adaptation problem is modeled as a MDP, where the total peak water level observed with certainty defines the states (*S*) and hence informs the actions (*A*) taken in time. In this work, the total peak water level incorporates the combination of local sea level rise (SLR), and storm surge (from hurricane events, for example). The following subsections briefly describe the models used to simulate the SLR and surge components respectively, along with the formulation of the MDP for flood management.

Sea Level Rise Model: The SLR model is derived from the SLR projections provided in the most recent assessment report (AR6) of the (IPCC) (Fox et.al. 2021). The regional SLR projections capture the long-term trends in sea levels due to changes in weather patterns and other factors related to climate change. Additionally, in order to capture the short-term variability present in the historical observations, a regression model is fit to the observed historical data provided by the National Oceanic and Atmospheric Administration (NOAA) (NOAA 2023), whereby the random noise present in the observations is assumed to be Gaussian with a fixed standard deviation. Sample SLR realizations are shown in Figure 4(a).

Storm Surge Model: Storm driven surges are simulated as annual maximum surge heights from a generalized extreme value (GEV) distribution. The parameters of the stationary GEV distribution are determined by maximum likelihood estimation on the observed annual maximum tidal gauge readings at the Battery (Ceres et.al. 2017), and result in an estimated 100-year surge

Downloaded from ascelibrary.org by Gordon Warn on 11/30/23. Copyright ASCE. For personal use only; all rights reserved.

level of 2.62 m above the MHHW (Mean Higher High Water) datum, consistent with other estimates available in the literature (Lin 2010).

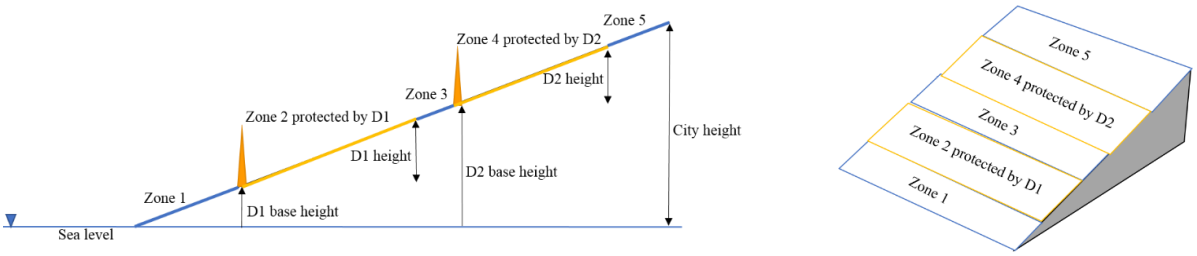

Figure 3: Coastal City inspired by Manhattan.

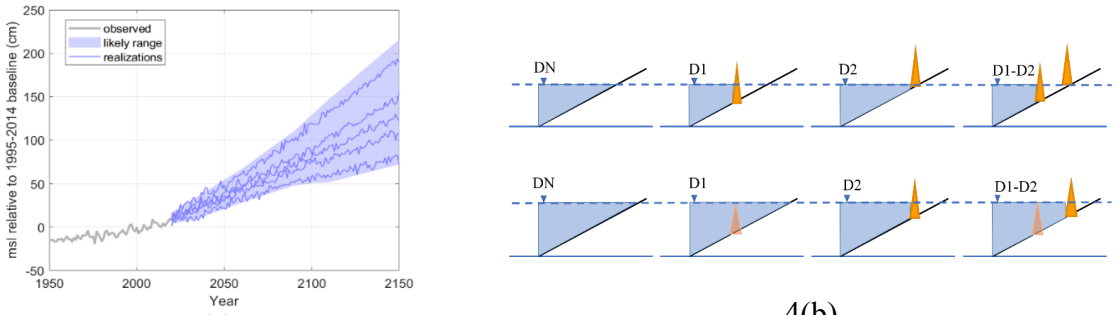

4(a) 4(b)

Figure 4. Illustration of sea level rise (SLR) process simulations and rewards model. SLR realizations (msl denotes mean sea level) based on IPCC SSP2-4.5 projections along with variability present in past observations is shown in 4(a). In 4(b), it can be seen that the flooded regions vary across the 4 considered systems DN (system with no floodwalls), D1 (system with only the lower floodwall), D2 (system with only the higher floodwall) and D1-D2 (system with both the floodwalls).

State Space: The state of the MDP environment in this work is the total peak water level, consisting of two components; a local sea level rise and a storm-driven surge height. To formulate a discrete-state MDP, both the continuous state variables are suitably discretized. The SLR variable is discretized with a step size of 2 cm with the maximum SLR value being 150 cm, resulting in a total number of 77 discrete SLR states $\left(S_{SLR}\right)$. Similarly, the storm surge variable is defined by 72 discrete surge levels $\left(S_{surge}\right)$, with a step size of 10 cm until the maximum surge level of 700 cm. The two components are combined to obtain a total of 5,544 states, as $S_t = S_{SLR} \times S_{surge}$. Since the sea level rise process is non-stationary, time or the SLR rate needs to be incorporated in the state space. A 40 years horizon has been considered in this work and, thus, the state space consists of 77 primary SLR states $\left(S_{SLR}\right)$, considered for each of the 40 different time points, resulting in a total of 3,080 SLR states $\left(S'_{SLR}\right)$, combined with 72 surge states $\left(S_{surge}\right)$, resulting in turn in a total of 221,760 states, as $S_t = S'_{SLR} \times S_{surge}$.

The available actions affect how the system will evolve in the future. For example, the construction of a dike at some point in time will change the region that is flooded under a future

Downloaded from ascelibrary.org by Gordon Warn on 11/30/23. Copyright ASCE. For personal use only; all rights reserved.

Downloaded from ascelibrary.org by Gordon Warn on 11/30/23. Copyright ASCE. For personal use only; all rights reserved.

storm of a given height, as shown in Figure 4(b), and hence damage, which is related to the rewards. To account for different rewards modeled under different configurations of the system, which are possible with the given set of actions, the state space needs to include the states under each possible system configuration which can be controlled. With two possible interventions, the two dikes D1 and D2 as shown in Figure 3, there can overall be 4 possible system configurations; DN where none of the measures have been implemented yet, D1 where only the lower dike D1 is constructed, D2 where only the higher dike D2 is constructed, and lastly the terminal system D1-D2 where both the dikes have been built, as shown in Figure 4(b). Thus, with two possible interventions, the state space is augmented here to incorporate 221,760 states belonging to each of the controllable systems. Finally, the total state space is obtained by adding a terminal state for a finite horizon formulation.

Actions: The actions $a \in A$ represent the state-dependent flood protection actions, such as construction of barriers, dikes, or seawalls, as well as green infrastructure options, like repurposed land, salt-marsh, and oyster reefs. As mentioned previously, there can be four system configurations with two possible flood protections and the effect of taking any action (apart from do-nothing action) is to make a transition from the current state s_t in system X, at time step t, to state s_{t+1} in system X', at the subsequent time step $t+1$, where X' represents the system containing the flood prevention measure associated with the action. Under do-nothing action, the environment continues in the same system configuration based on the transition probabilities derived from the climate projection models.

Transition Probabilities: The transition probability matrix $\mathbf{T}(s' \mid s, a)$ models the probability $P(s' \mid s, a)$ of transitioning to state $s' \in S$ from state $s \in S$ after taking an action $a \in A$. The transition probabilities are derived from the non-stationary SLR process and the storm surge model. With sufficient realizations of the SLR process, the SLR state transition probability $P_{SLR}^{t \to t+1}\left(s' \mid s, a_{DN}\right)$ under the do-nothing action for time t is calculated by counting the number of transitions from $S_t = s$ to $S_{t+1} = s'$ divided by the total number of observations in state s at time t. Hence, the 77 × 77 SLR transition matrix for each year in the future has been estimated by 100 million Monte Carlo simulations of the SLR process, a few of which are shown in Figure 4(a). The stationary storm surge transition probabilities $P_{surge}\left(s' \mid s, a\right)$ are straightforwardly obtained from the distribution of the underlying generalized extreme value (GEV) model. Subsequently, the transition probabilities for the combined states $P_{total}^{t \to t+1}\left(s' \mid s, a\right)$ are obtained by suitably taking the combination of each surge and SLR state, and calculating their transition probabilities based on their individual transition matrices. After obtaining the 5,544 by 5,544 transition probability matrix $\mathbf{T}^{t \to t+1}(s' \mid s, a)$ for each t of the planning horizon, the complete transition matrix is constructed as follows:

$$\mathbf{T}_{(221,760 \times 221,760)} = \begin{bmatrix} \ddots & \vdots & \iddots \\ \cdots & \mathbf{T}(i \to j) & \cdots \\ \iddots & \vdots & \ddots \end{bmatrix}; \mathbf{T}(i \to j)_{(11,088 \times 11,088)} = \begin{bmatrix} \mathbf{0}^{ii}_{(5,544 \times 5,544)} & \mathbf{T}^{ij}_{(5,544 \times 5,544)} \\ \mathbf{0}^{ji}_{(5,544 \times 5,544)} & \mathbf{0}^{jj}_{(5,544 \times 5,544)} \end{bmatrix} \tag{3}$$

where $\mathbf{T}_{(221,760 \times 221,760)}$ as shown in Eq. (3) is further developed considering the different system configurations and each action's effects on the system. Thus, the complete transition matrix

associated with each action is constructed by stacking the $\mathbf{T}$ in Eq. (3) according to the deterministic transition between the systems under the effect of the action.

Rewards: The reward $R(s,a)$ is defined as a function of the state s and action a. In the illustrative example, the following cost categories are considered: (i) expected flood damage costs (or risks), (ii) implementation/construction costs, and (iii) maintenance costs of the flood protection measures. The maintenance of any flood protection asset is conducted every year after the asset is constructed. The societal and environmental impacts, modeled as greenhouse gas (GHG) emissions associated with each monetary cost category are also incorporated. The life-cycle GHG emissions are translated into a net present monetary value using the Social Cost of Carbon (SCC), defined as the net present value of the long-term climate change impacts caused per ton of carbon dioxide equivalent released into the atmosphere at a certain point in time (Nordhaus 2017).

For the do-nothing (a_{DN}) action with no protective measure being implemented, the reward $R(s,a_{DN})$ corresponds to the costs associated with flood damages $C_f(s')$ that can be expected from transitioning to any state s' from the current state s. Since the environment continues in the same system for the DN action, the flood risk is calculated under the current system. Also, since maintenance of an asset is assumed to be performed every year after the asset is constructed, the maintenance cost (C_m) is also incorporated in the costs for the DN as:

$$R(s_X,a_{DN}) = C_f(s_X,a_{DN}) + C_m(s_X) = \sum_{s'\in S}\left[P(s'_X \mid s_X,a_{DN})C_f(s'_X)\right] + C_m(s_X) \qquad (4)$$

where s_X and s'_X represent the water level states s and s' respectively, both in system X.

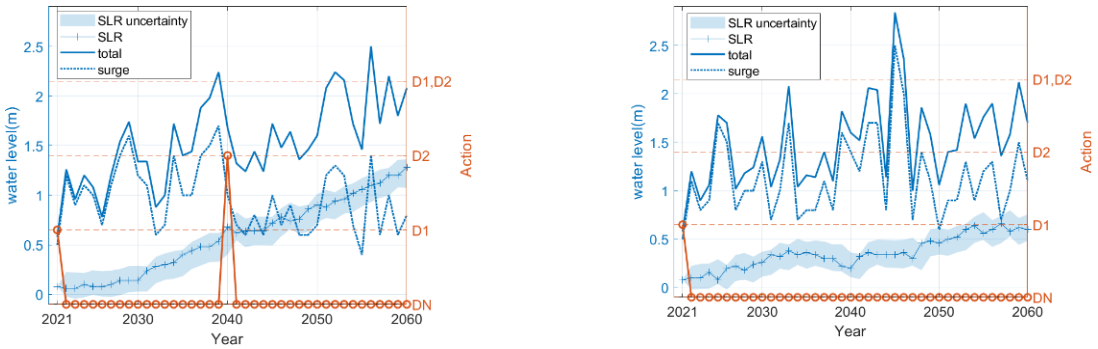

5(a) 5(b)

Figure 5. Policy realizations with two dikes scenario. The second dike is constructed only when the observed sea level rise (SLR) process progresses more rapidly. The shaded region associated with the SLR process represents the uncertainty over the next time step from the current time step.

For any other action apart from DN, the cost includes the expected flood damage costs along with the implementation cost (C_i) of the flood mitigation action (a_i). Since, any action apart from DN results in a transition from a system to another, the flood risk is calculated according to the newly changed system as:

Downloaded from ascelibrary.org by Gordon Warn on 11/30/23. Copyright ASCE. For personal use only; all rights reserved.

$$R(s_X, a_i) = C_f(s_X, a_i) + C_i(a_i) = \sum_{s' \in S} \left[P(s'_{X'} \mid s_X, a_i) C_f(s'_{X'}) \right] + C_i(a_i) \quad (5)$$

where s_X and $s'_{X'}$ represent the water level states s and s' in systems X and X' respectively.

CASE STUDY: RESULTS AND DISCUSSSIONS

For the purpose of this study, only a representative unit meter width of the idealized coastal city, as shown in Figure 3, parallel to the coastline is considered. Due to the exponentially increasing state space with the number of possible interventions, as previously explained, the policies presented here consist of two dynamic actions which can be implemented adaptively in time. For the example shown here, the two dynamic actions considered are two possible dikes D1 and D2, as shown in Figure 3 at two different fixed locations. It is to be noted that scaling up of the state and action spaces can be achieved with deep reinforcement learning based solution methods (Ororbia and Warn 2023, Andriotis and Papakonstantinou 2021), albeit at the cost of loss of global optimality guarantee.

Realizations of computed MDP policies are shown in Figure 5. With the SLR process being a non-stationary evolving process, the sequential decisions to be taken are hence primarily driven by the SLR. Thus, in the scenario with two dikes, the higher dike is erected only when an extreme level of SLR is observed. Otherwise, the first dike at a lower elevation is constructed at the beginning of the management horizon, illustrating the urgency of flood risk mitigation actions. The immediate action of constructing the lower dike can be understood by the high vulnerability of the city, with the highest possible flood damage being $ 1 million (per m of shoreline), when flood height reaches the top height of the city at 8.5m.

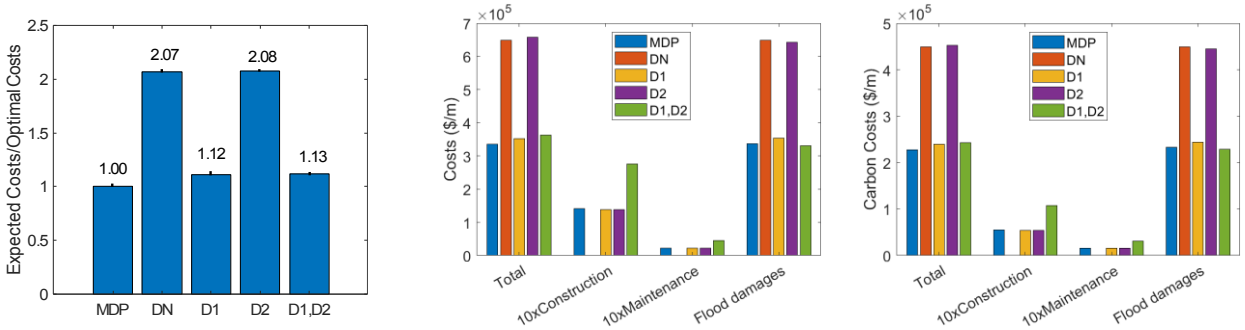

6(a) 6(b) 6(c)

Figure 6. Overall expected costs are shown in 6(a) including all components of costs and carbon costs. Costs and carbon costs decomposition shown in 6(b) and 6(c), respectively. Total (construction + maintenance + flood damage), construction, maintenance and flood damage expected costs are in $ (construction and maintenance costs are amplified 10 times for visual purposes). The static baselines are: DN-no measures taken, D1- only dike D1 constructed, D2- only dike D2 constructed, and D1, D2- both dikes constructed at the beginning of planning horizon.

In comparison to static benchmarks, shown in Figure 6, the total overall (monetary and carbon) costs of the MDP policy are the lowest, resulting in a reduction of over $0.09 million

Downloaded from ascelibrary.org by Gordon Warn on 11/30/23. Copyright ASCE. For personal use only; all rights reserved.

compared to the closest static solution. This empirical result agrees with the theoretical convergence guarantees of dynamic programming solutions to global optimality. The carbon cost associated with any cost category shown here is calculated as the sum of social carbon costs accumulated over the future horizon, from the time of carbon emission in the atmosphere. Because of carbon accumulation and the increasing social cost of carbon in the future (NYSDEC 2022), the carbon costs become comparable to the monetary costs and hence can play a significant role in driving adaptation policies.

CONCLUSIONS

As part of ongoing work, in this paper flood risk management under uncertainties posed by climate change has been framed as a stochastic-control process, to identify the optimal flood protection measures to be taken adaptively over time. The efficacy of the adaptive framework is demonstrated by the reduction in the overall costs of flood management in comparison to those associated with static methods, due to better-informed policies which are chosen according to the actual evolving climate conditions, through a closed-loop format that guarantees global optimality. Overall, the presented work aims to introduce a new paradigm of life-long optimal infrastructure adaptation based on observed data and realized conditions, and responding to the important uncertainties that characterize the problem. One challenge of this presented framework is the growing computational complexity with an increase in the number of possible dynamic actions and length of the time horizon, which cannot be handled with MDP solvers with guarantees of global optimality. To overcome this, further research is suggested in the direction of deep reinforcement learning based solution techniques.

ACKOWLEDGEMENTS

The authors would like to acknowledge the support of the U.S. National Science Foundation (NSF), which supported this research under Grant No. CMMI-2053620

REFERENCES

Andriotis, C. P., and Papakonstantinou, K. G. (2021). “Deep reinforcement learning driven inspection and maintenance planning under incomplete information and constraints.” *Reliability Engineering & System Safety*, 212, 107551.

Arias, P., et al. (2021). Climate change 2021: The physical science basis. *Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change; Technical Summary*.

Ceres, R. L., Forest, C. E., and Keller, K. (2017). Understanding the detectability of potential changes to the 100-year peak storm surge. *Climatic Change*, 145, 221-235.

Ceres, R. L., Forest, C. E., and Keller, K. (2019). “Optimization of multiple storm surge risk mitigation strategies for an island city on a wedge.” *Environmental Modelling & Software*,119, 341-353.

Fox-Kemper, B., et al. (2021). Ocean, Cryosphere and Sea Level Change. In: Climate Change 2021: The Physical Science Basis. *Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*.

Downloaded from ascelibrary.org by Gordon Warn on 11/30/23. Copyright ASCE. For personal use only; all rights reserved.

Cambridge University Press. (Forthcoming). Sea Level Projections from the IPCC 6th Assessment Report. <https://podaac.jpl.nasa.gov/announcements/2021-08-09-Sea-level-projections-from-the-IPCC-6th-Assessment-Report>(accessed June 20 2023).

Lin, N., Emanuel, K. A., Smith, J. A., and Vanmarcke, E. (2010). Risk assessment of hurricane storm surge for New York City. *Journal of Geophysical Research: Atmospheres*, 115.

New York State Department of Environmental Conservation. (2022). Appendix: NYS Social Cost Values. <https://www.dec.ny.gov/docs/administration_pdf/vocapp22.pdf>(accessed June 20 2023).

NOAA. (2023). NOAA Tides & Currents, <https://tidesandcurrents.noaa.gov/waterlevels.html?id=8518750>(accessed June 20 2023).

Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114, 1518–1523.

Ororbia, M. E., and Warn, G. P. (2023). "Design synthesis of structural systems as a Markov decision process solved with deep reinforcement learning." *Journal of Mechanical Design*, 145(6), 061701.

Puterman, M. L. (2014). *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons.

Papakonstantinou, K. G., and Shinozuka, M. (2014). "Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part I: Theory." *Reliability Engineering & System Safety*, 130, 202-213.

Smith, T., and Simmons, R. (2006). Focused real-time dynamic programming for mdps: Squeezing more out of a heuristic. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1227-1232.

Ward, P. J., et al. (2017). "A global framework for future costs and benefits of river-flood protection in urban areas." *Nature Climate Change*, 7, 642–646.

Downloaded from ascelibrary.org by Gordon Warn on 11/30/23. Copyright ASCE. For personal use only; all rights reserved.