14th International Conference on Structural Safety and Reliability - ICOSSAR'25
1-6 June, 2025 - Los Angeles California, USA
R. Ghanem, E. Johnson, S. Masri and A. Olivier (Eds)

OPTIMAL ADAPTIVE INFRASTRUCTURE PLANNING UNDER CLIMATE CHANGE

ASHMITA BHATTACHARYA, ADITYA SHARMA, GORDON P. WARN and KOSTAS G. PAPAKONSTANTINOU*

The Pennsylvania State University
University Park, PA, 16802, United States

Key words: stochastic control, climate change, sea level rise, nature-based infrastructure, carbon emissions, system effects

Abstract. Various climate change effects pose increasing risks to our built environment and infrastructure. Existing approaches for managing climate-related implications typically employ a risk-based, cost-benefit analysis that evaluates a comprehensive set of mitigation strategies against a wide range of simulated possible future scenarios. However, the problem is characterized by substantial aleatoric as well as epistemic uncertainties, and cost-benefit strategies can often lead to policies that might be optimal in an average sense, over the mean of anticipated future scenarios, but cannot offer optimal adaptive solutions based on the actual climate effects that eventually evolve in time. To address this challenge, we instead formulate climate risk management and related adaptive infrastructure planning as a formal dynamic framework for sequential decision-making. In particular, our methodology follows a closed-loop stochastic control-based approach using Markov Decision Processes (MDP) and Partially Observable MDPs (POMDPs), taking real-time data into account for evaluating the evolving conditions, selecting the best possible life-cycle actions in time. The developed framework is mainly illustrated here through two coastal adaptation applications. Both applications consider the storm surge and sea-level rise hazards as well as a set of diverse action types, including grey- as well as nature-based solutions. Environmental impacts, such as carbon emissions and sequestration, via the notion of the social cost of carbon, are also accounted for, offering a holistic approach that addresses both the economic and environmental dimensions of coastal flood protection.

1 INTRODUCTION

Climate change is impacting every region globally, with rising greenhouse gas emissions causing global temperatures to increase. This results in sea level rise (SLR) through land ice loss and ocean warming [1], posing severe risks to coastal communities that are projected to worsen by the century's end. Effective flood protection policies are urgently needed, yet optimizing them is challenging due to constrained resources and multiple layers of uncertainty [2, 3]. These uncertainties include natural variability in climate processes, imperfect knowledge of climate models, potential updates to these models, future trajectory evolution under assumed scenarios, and time-dependent observational limitations. While traditional planning, particularly among federal agencies, has relied on static cost-benefit analysis (CBA) [4, 5], static optimization frameworks often yield a single policy optimized for average or percentile-based climate projections. Such approaches, focused on methods like genetic algorithms [6], decision trees [2], and non-linear programming [7], fail to incorporate real-time data, leaving policies vulnerable to sub-optimality as actual climate trajectories deviate from assumed projection means (or chosen percentiles) [1], as shown in Figure 1. Dynamic approaches leveraging evolving information are thus critical for refining risk assessments and addressing the associated uncertainties effectively.

Adaptive planning approaches that adjust policy actions based on evolving climate conditions have gained attention in the literature, employing methods such as direct policy search (DPS) [9, 10], real options analysis (ROA) [11, 12], and stochastic control [13, 14, 15]. These methods have primarily been applied to levee construction and heightening scenarios, where a single levee is incrementally modified over time. More recent studies leverage stochastic control frameworks, formulating decision-making as Markov Decision Processes (MDPs) or Partially Observable MDPs (POMDPs), with state variables representing conditions and decisions corresponding to infrastructure actions. These approaches can optimize expected costs or rewards over a decision horizon by determining actions conditioned on observed or inferred states, offering global optimality guarantees, within the framework settings, through dynamic programming and closed-loop stochastic control. Notable works, however, still primarily focus on levee heightening applications [13, 14, 16], not demonstrating the feasible broader applicability to scenarios involving diverse actions and/or more complex systems.

This work introduces significant advancements in adaptive planning and a solution framework for optimal life-cycle infrastructure adaptation under climate change. Leveraging the mathematical foundations of MDPs/POMDPs, we provide a general, adaptable, and extensible approach with higher-fidelity representations compared to existing methods. Our framework integrates current IPCC models [1], accommodates diverse action types (e.g., floodwalls, nature-based infrastructure, seawalls), addresses model uncertainties, and incorporates the social cost of carbon. As shown in Figure 1B, integrating annual environmental data significantly reduces uncertainty.

The proposed framework applies MDPs to optimize coastal infrastructure adaptation under storm surge and sea level rise (SLR) risks across two distinct applications. An MDP, defined as a 5-element tuple $\langle S, A, P, R, \gamma \rangle$, models a controlled stochastic process where the system state $s \in S$ transitions probabilistically to s' based on Markovian dynamics $P(s'|s, a)$ after action $a \in A$, yielding a reward or cost $R(s, a)$. In this framework, the states are defined using discretized SLR states, storm surge

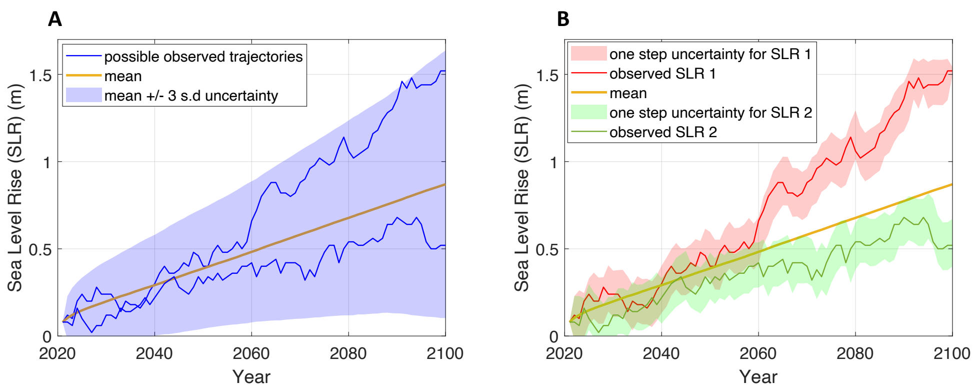

Figure 1: The representative uncertainty involved in decision-making. Figure 1A illustrates the uncertainty range of the IPCC AR6 climate projection model (SSP5-8.5) from the current time step, while Figure 1B depicts uncertainty bands for sequential decision-making at each step [8].

levels, and infrastructure conditions. The SLR states evolve probabilistically according to IPCC AR6 projections [1], with the infrastructure states shaped by the history of past adaptation actions. Unlike basic levee heightening policies, our framework accounts for the cumulative effects of prior actions. In this work, related costs at each time-step consist of flooding risks, maintenance of pre-existing flood protection assets, and implementation costs of new flood protection measures associated with the action being taken. Thus, the optimal value function V^*, representing the expected γ-discounted rewards under the optimal policy π^* from a given system state s, is given as:

$$V^*(s) = \max_{a \in A} \left[\left\{ R_i(s,a) + R_m(s,a) + \gamma \sum_{s' \in S} p(s'|s,a) R_f(s') \right\} + \gamma \sum_{s' \in S} p(s'|s,a) V^*(s') \right] \quad (1)$$

where R_i, R_m, and R_f are costs associated with the implementation of the new flood measures, maintenance of existing measures, and flood-induced damages, respectively. Carbon emissions associated with each one of these components and related implications are also considered in our framework.

In this work, we also extend our framework to a POMDP formulation, where hidden states correspond to the underlying climate models, represented by the Shared Socio-economic Pathways (SSPs), inferred probabilistically from observed SLR levels. The belief state, $\mathbf{b}$, updated using Bayes' rule, encapsulates a probability distribution over SSPs, serving as a sufficient statistic for past observations and guiding optimal decision-making [17].

The environmental impacts of coastal flood protection are also often overlooked. While studies highlight the economic costs and benefits of protection [18, 19], few quantify the life-cycle environmental impacts [20]. This work also fills this gap by introducing a dynamic framework that incorporates life-cycle greenhouse gas (GHG) emissions, balancing economic and environmental factors in flood risk mitigation.

2 SEA LEVEL RISE AND STORM SURGE MODELS

In this work, the MDP state variables are defined as the possible total peak water levels that can be realized in time, which incorporate the combination of local sea level rise and storm surge (from hurricane events, for example). The SLR model leverages CMIP6 projections from IPCC AR6 [1] for the Battery tidal gauge in NYC, extending through 2150, as shown in Figure 2A. These projections consider various climate scenarios based on SSPs and utilizes 21 Global Climate Models (GCMs) per SSP to capture epistemic uncertainties effectively. Projections follow a chosen percentile within the likely range, accounting for long-term trends from thermal expansion and ice sheet melt, while short-term variability is modeled using a regression of NOAA historical data [21] with Gaussian noise, as shown in Figure 2B. Simulations combining IPCC projections and historical variability are shown in Figure 2C. Storm-driven surges are modeled as annual maximum surge heights using a generalized extreme value (GEV) distribution, with parameters estimated from Battery tidal gauge data after removing local SLR trends, resulting in a 100-year surge level of 2.62 m above the Mean Higher High Water (MHHW) datum, consistent with literature [22], as shown in Figure 2D.

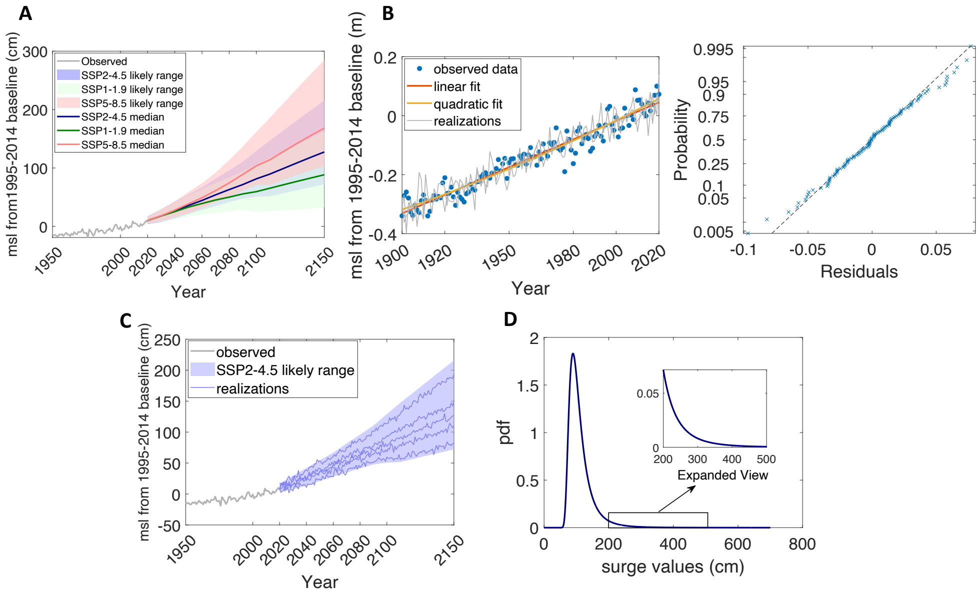

Figure 2: (A) Regional sea level change for NYC, with likely range (5th-95th percentiles). (B) Historical variability, fitted polynomials, and Gaussian-distributed hindcast samples. (C) SLR realizations from IPCC SSP2-4.5 projections with historical variability. (D) GEV probability density function for storm surges [8].

3 RESULTS

This paper examines adaptive flood management in two distinct coastal applications. Leveraging our MDP framework, annual peak water levels are used as state variables to inform actions, while future SLR is mainly modeled under the SSP 2-4.5 scenario.

3.1 Application 1: Adaptive Planning for Individual Coastal Regions

Here, we explore two coastal settings: a coastal city (inspired by Manhattan) and a coastal community (inspired by Staten Island).

3.1.1 Coastal City Setting

This study focuses on decision-making for flood risk mitigation in a coastal city with no pre-existing buffers, where fixed-height floodwalls are erected. Sample MDP policies (Figure 3A) demonstrate the urgency of adaptive flood risk mitigation. The lower-elevation floodwall is constructed early, reflecting the city's vulnerability, with potential flood damages reaching \$1 million per meter of shoreline at an 8.5m flood height, while the higher-elevation floodwall is built only under extreme SLR conditions. The absence of existing protective assets, such as seawalls or wetlands, underscores the need for proactive interventions. Monte Carlo simulations (1,000 runs) validate the effectiveness of the computed MDP policies, achieving lower total costs compared to static baselines (Figure 3B), with savings

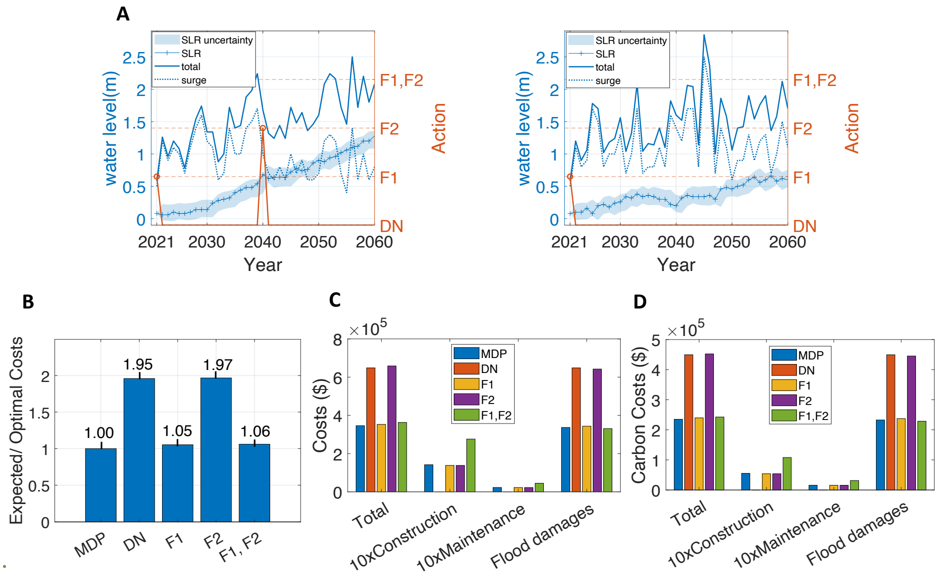

Figure 3: (A) Two policy realizations with shaded uncertainty bands associated with the SLR process: the second floodwall is built only under rapid SLR progression. (B) Expected normalized costs of MDP-based policy compared to static baselines (DN-no measures taken, F1- only floodwall F1 constructed, F2- only floodwall F2 constructed, and F1, F2- both floodwalls constructed at the beginning of planning horizon). (C) and (D) show breakdowns of monetary ($) and carbon costs, including construction, maintenance, and flood damage (construction/maintenance costs amplified for clarity) [8].

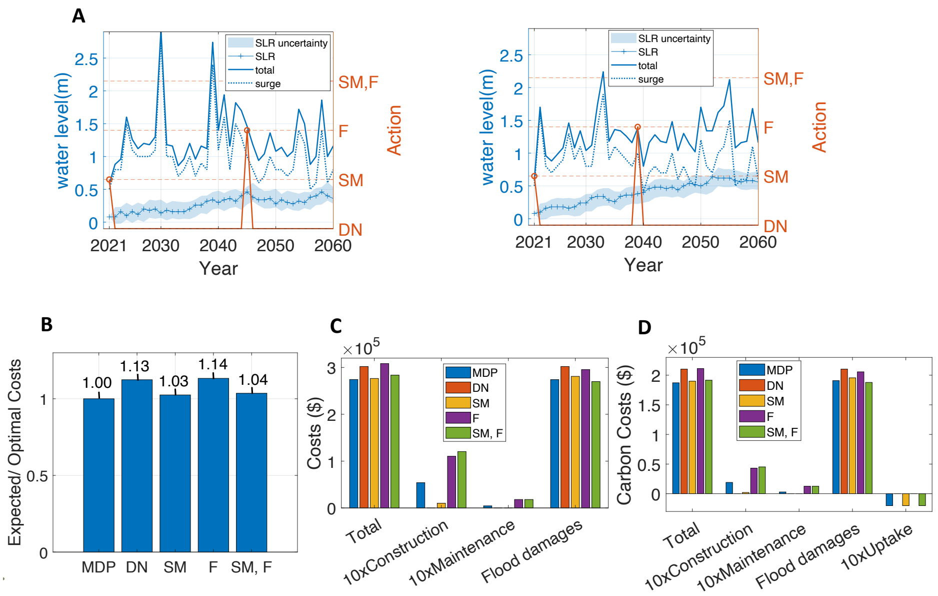

Figure 4: (A) shows two policy realizations where the salt marsh is constructed immediately, followed by adaptive construction of the higher floodwall. (B) Expected normalized costs of MDP-based policy versus static baselines (DN-no measures taken, SM- only salt marsh constructed, F- only floodwall F constructed, and SM, F- both measures constructed at the beginning of the planning horizon). (C) Monetary costs (construction, maintenance, flood damage) and (D) carbon costs, including uptake rewards, in $ [8].

exceeding \$0.03 million per meter of coastline. Cost breakdowns (Figure 3C) show monetary expenses across construction, maintenance, and flood damages, while carbon costs (Figure 3D) reveal the significant role of carbon emissions in decision-making. The cumulative carbon costs, influenced by future increases in the social cost of carbon [23], approach monetary costs, highlighting their substantial impact on driving adaptive policies. We also explore a scenario where shoreline land is repurposed into a multifunctional green zone, integrated with an existing seawall, to provide cost savings and carbon sequestration benefits. In this case, adaptive policies consistently prioritize the timely development of the green zone in response to rising water levels, eliminating the need for the higher-elevation floodwall. Representative policies and additional details are provided in [8].

3.1.2 Coastal Community Setting

This coastal low-lying area, inspired by Staten Island, NY features a flatter terrain than the Manhattan-inspired urban setting. While more flood-prone, the resulting damages are lower. In this coastal community setting, dynamic actions include constructing a fixed-height floodwall (1.5 m) or selecting nature-based infrastructure, such as salt marshes or oyster reefs. These green infrastructure options provide wave attenuation along with some carbon sequestration benefits.

Figure 4A illustrates sample MDP policy realizations for the salt marsh scenario, with similar policies observed for the oyster reef scenario as detailed in [8, 24]. In both cases, the floodwall is constructed when sea levels approach 0.5 m. Normalized total policy costs for the salt marsh scenario are depicted in Figure 4B, with detailed monetary costs in Figure 4C and carbon costs in Figure 4D. The oyster reef scenario results, available in [8], reveal lower overall costs due to superior flood attenuation, though with negligible carbon uptake compared to the salt marsh. However, the salt marsh's carbon uptake is insufficient to offset higher flood damage caused by its reduced wave attenuation. MDP-based policies consistently yield lower overall costs, achieving savings of approximately \$0.018 million per meter of coastline for the salt marsh scenario and \$0.01 million for the oyster reef scenario, relative to static baselines. These findings confirm that MDP-based solutions effectively converge to globally optimal outcomes, even in complex scenarios with closely competing local solutions.

3.1.3 Effect of Social Cost of Carbon and Climate Model

The social cost of carbon quantifies the economic impact of greenhouse gas emissions on society, accounting for both present and future climate change damages [25, 26]. The Environmental Protection Agency (EPA) has proposed raising this estimate to \$190 per ton, nearly four times higher than the current value of \$51 per ton [27]. In this work, the effect of a higher social cost of carbon on adaptation policies in the coastal city setting with two floodwalls is analyzed. An increase in the social cost of carbon leads to a higher frequency of constructing the floodwall at the higher elevation, as shown in Figure 5A, and encourages earlier interventions, as reflected in Figure 5B.

To study the effect of the underlying climate model on computed policies, two policies were trained on different climate model scenarios—one based on SSP2-4.5 and the other on the more intense SSP5-8.5 scenario. Since both policies are driven by real-time observations, Figure 5C shows that they are similar when climate trajectories align, demonstrating that the MDP-based framework can produce similar adaptive actions under similar conditions, even with different models. However, the models only begin to significantly diverge after 2100. To incorporate climate model uncertainty, the framework was extended to a POMDP-based formulation, where SSP2-4.5 and SSP5-8.5 are considered hidden states. Bayesian principles are used to update beliefs about the models based on observed SLR, as shown in Figures 5D-F. For a given SLR trajectory, the belief evolution indicates how different SLR patterns affect the likelihood of each model. The POMDP policies show similar behavior to MDP policies, with higher floodwalls constructed under higher SLR scenarios. While the models' effects did not significantly influence the derived policy in this study, a more detailed future analysis could assess how diverse climate outputs affect MDP/POMDP decision-making performance.

3.2 Application 2: Adaptive Planning for Interconnected Coastal Regions

While Section 3.1 demonstrated the effectiveness of adaptive decision-making considering individual regions, infrastructure systems in coastal environments can also be inherently interconnected,

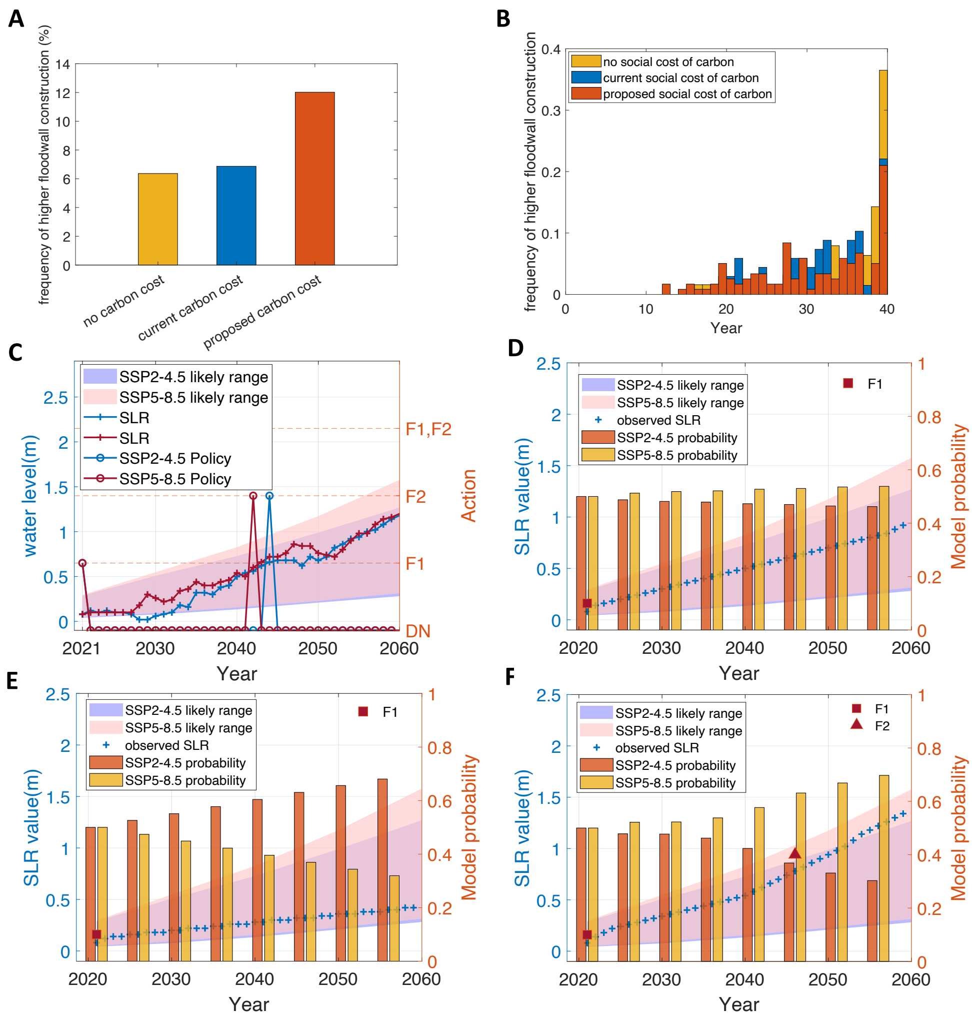

Figure 5: (A) Impact of social cost of carbon (SCC) on floodwall construction frequency and (B) timing: no SCC, current ($51/ton), and proposed ($190/ton) values. (C) Policies trained on different SSPs exhibit similar behavior under matching climate observations. (D-F) Belief evolution over SSP scenarios every 5 years, based on three observed SLR trajectories, with actions F1 and F2 highlighted. F1 and F2 in the legends of Figures (D), (E) and (F) correspond to the actions F1 and F2 being taken, respectively [8].

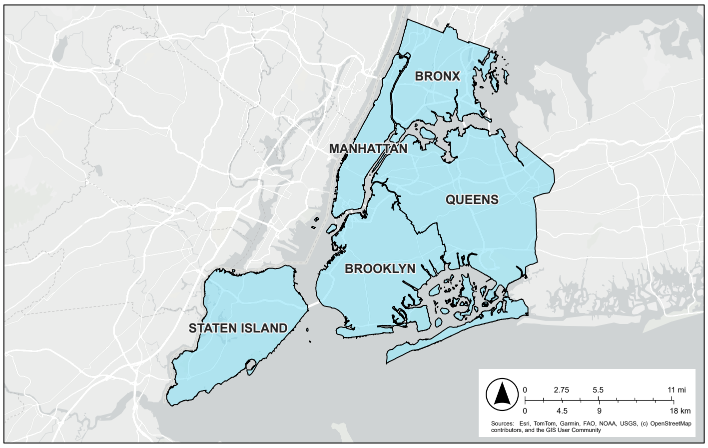

Figure 6: NYC coastal environment.

forming complex networks where the flooding in one area may influence, or be influenced by, conditions in neighboring regions. In such interdependent infrastructure systems, it is essential to consider coordinated and system-wide optimal adaptation policies, since ignoring system effects may lead to underestimation of cascading flood risks and policies that are optimal only in isolation.

To investigate and demonstrate the effectiveness of adaptive flood management considering interconnected systems, this application offers a first, conceptual modeling of the coastal environment of New York City (NYC) as a system composed of its five boroughs, following their topographical features and economic exposure. Adaptation actions across the city characterize a diverse set of feasible strategies that reflect local shoreline typologies, land-use constraints, and protection needs across the different boroughs. For the Bronx and Brooklyn, two elevation-tiered floodwalls (F1 and F2) are considered, enabling staged responses to observed water levels. Queens features the option of deploying oyster reefs (OR) to attenuate incoming waves and a higher elevation floodwall (F2). Staten Island, characterized by its low elevation and open coastal setting, considers the construction of a salt marsh (SM) and a higher floodwall (F2) as viable protection measures. For Manhattan, which is already protected by a 1.2 meter seawall, a multi-functional green resistance (GR) zone and a higher-elevation floodwall (F2) are considered. System-level interactions are captured by modeling flooding interactions between relevant boroughs depending upon spatial proximity, shared infrastructure, and relative elevation differences, as shown in Figure 6.

For this preliminary application, two sample realizations of the computed MDP policies are shown in Figure 7A. As seen, the construction of higher floodwalls is dependent on the observed SLR trajectories. Although the vulnerabilities of the different boroughs are quite different, the critical floodwall (F2) for all system components is constructed within a similar time frame to prevent lateral flooding effects. Given the high vulnerabilities and low-lying nature of Staten Island, Brooklyn, and Queens, the corresponding adaptation actions available at the lower elevations (SM, F1, and OR) are constructed at the beginning of the planning horizon. For Manhattan, the multi-purpose green resistance is optimally constructed in time according to the evolving condition.

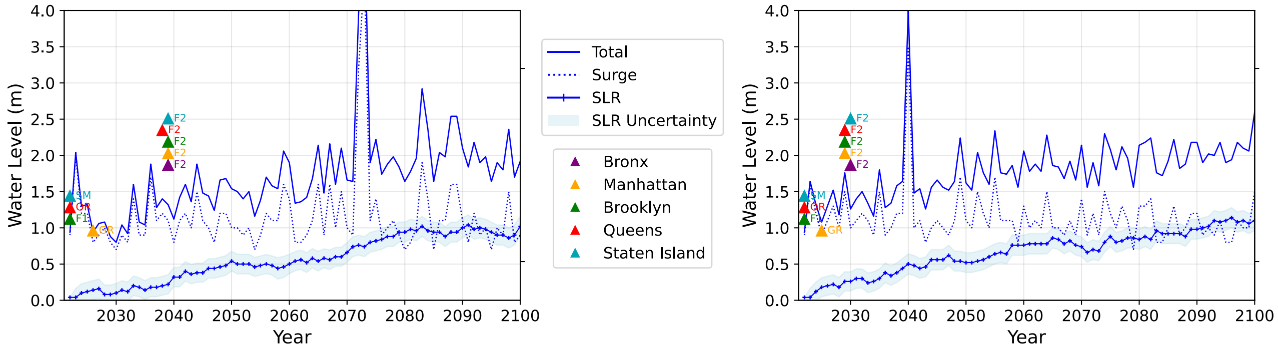

(A)

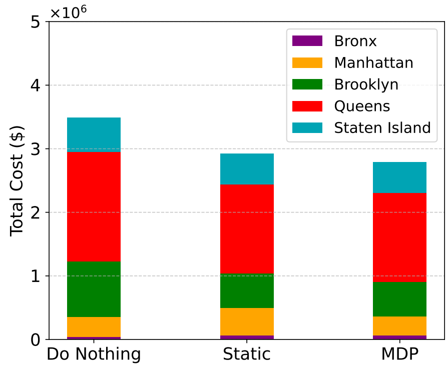

(B)

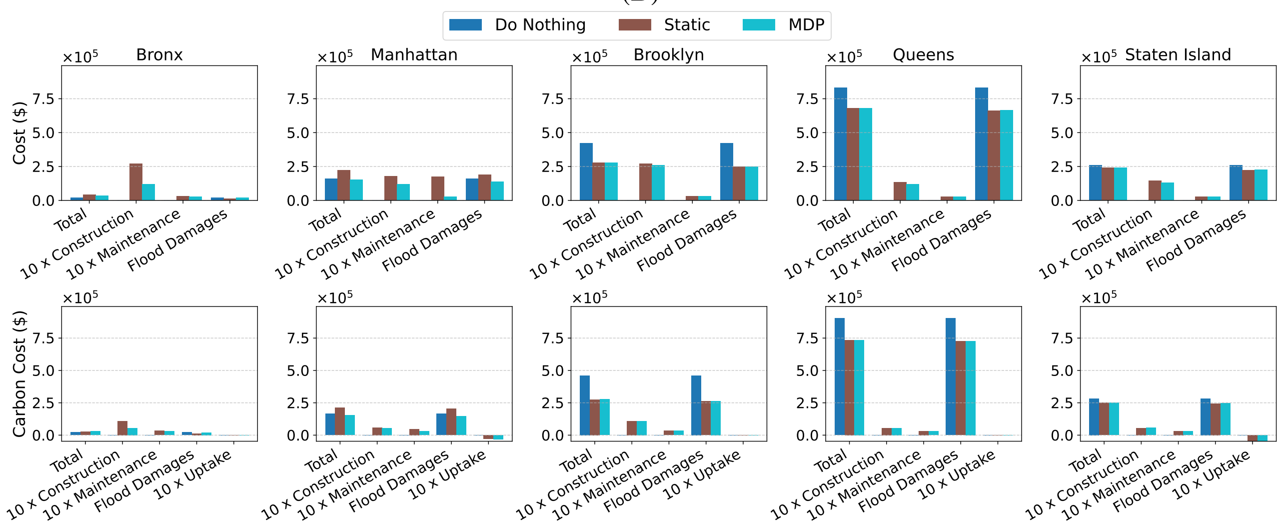

(C)

Figure 7: (A) Policy realizations for the NYC coastal system demonstrating adaptive actions at different times in different boroughs depending upon the observed water levels. (B) Total system costs under different planning strategies. (C) Decomposition of individual cost components, including construction, maintenance, flood damages, and carbon costs across all five boroughs.

The overall expected system cost achieved by the MDP policy for this application is compared in Figure 7B with the corresponding static, cost-benefit baseline. By dynamically adjusting to observed conditions, the MDP policy achieves the lowest total cost, saving over $\$0.134$ million per unit length of coastline compared to the closest static solution. The decomposition for monetary and carbon costs, along with the corresponding cost components for all five boroughs, is shown in Figure 7C. As seen, the cumulative carbon costs become comparable to the corresponding monetary costs, substantially impacting the policy decisions.

4 CONCLUSIONS

This work presents a stochastic control MDP/POMDP framework for life-cycle infrastructure adaptation under climate change, incorporating a variety of action types, including grey- as well as nature-based solutions. Using the concept of the social cost of carbon (SCC), environmental effects such as carbon emissions and sequestration are also taken into account, providing a comprehensive strategy that takes into account not just the economic, but also the environmental aspects of coastal flood protection. In addition to emphasizing the value of adaptive planning in time, the influence of the magnitude of the social cost of carbon and the underlying climate model assumptions on adaptive policies are also demonstrated. Currently, ongoing extensions of this framework to consider inherently interconnected coastal environments are also further explored, using NYC as an example, and considering flooding interactions between relevant regions and system effects.

While the presented MDP/POMDP-based approaches offer optimality guarantees, computational complexities grow with increasing states and time horizons. Future work will thus further explore multi-agent deep reinforcement learning solutions [28] and multi-objective formulations, and will compare solutions with other adaptive approaches in the literature. The growing importance of carbon costs in decision-making and the co-benefits of nature-based infrastructure for flood management and carbon sequestration are also worth further exploration.

5 ACKNOWLEDGMENTS

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

[1] Arias, Paola and Bellouin, Nicolas and Coppola, Erika and Jones, Richard and Krinner, Gerhard and Marotzke, Jochem and Naik, Vaishali and Palmer, Matthew and Plattner, G-K and Rogelj, Joeri and others, “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,” 2021.

[2] R. L. Sriver, R. J. Lempert, P. Wikman-Svahn, and K. Keller, “Characterizing uncertain sea-level rise projections to support investment decisions,” *Plos One*, vol. 13, no. 2, p. e0190641, 2018.

[3] D. M. Frangopol, M. Liu, M. Akiyama, D. Y. Yang, K. G. Papakonstantinou, K. Haas, M. G. Stewart, F. Biondini, M. Ghosn, S. Bianchi, *et al.*, “Life-cycle risk-based decision-making in a changing climate,” in *Effects of Climate Change on Life-Cycle Performance of Structures and Infrastructure Systems: Safety, Reliability, and Risk*, pp. 207–292, American Society of Civil Engineers (ASCE), 2024.

[4] Federal Emergency Management Agency (FEMA), "Guidelines for implementing executive order 11988 floodplain management and executive order 13690 establishing a federal flood risk management standard and a process for further soliciting and considering stakeholder input," 2015.

[5] E. Penning-Rowsell and D. Parker, "Economics of prevention measures addressing coastal hazards: Guideline 1 - implementation of cost-benefit analysis of coastal risk management prevention measures against natural hazards," in *European Commission, DG-ECHO, Civil Protection Unit*, 2016.

[6] R. L. Ceres, C. E. Forest, and K. Keller, "Optimization of multiple storm surge risk mitigation strategies for an island city on a wedge," *Environmental Modelling & Software*, vol. 119, pp. 341–353, 2019.

[7] C. Eijgenraam, J. Kind, C. Bak, R. Brekelmans, D. den Hertog, M. Duits, K. Roos, P. Vermeer, and W. Kuijken, "Economically efficient standards to protect the Netherlands against flooding," *Interfaces*, vol. 44, no. 1, pp. 7–21, 2014.

[8] A. Bhattacharya, K. G. Papakonstantinou, G. P. Warn, L. McPhillips, M. M. Bilec, C. E. Forest, R. Hasan, and D. Chavda, "Optimal life-cycle adaptation of coastal infrastructure under climate change," *Nature Communications*, vol. 16, no. 1, p. 1076, 2025.

[9] J. D. Quinn, P. M. Reed, and K. Keller, "Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points," *Environmental modelling & software*, vol. 92, pp. 125–141, 2017.

[10] G. G. Garner and K. Keller, "Using direct policy search to identify robust strategies in adapting to uncertain sea-level rise and storm surge," *Environmental Modelling & Software*, vol. 107, pp. 96–104, 2018.

[11] M. Woodward, Z. Kapelan, and B. Gouldby, "Adaptive flood risk management under climate change uncertainty using real options and optimization," *Risk Analysis*, vol. 34, no. 1, pp. 75–92, 2014.

[12] A. Wreford, R. Dittrich, and T. D. van Der Pol, "The added value of real options analysis for climate change adaptation," *Wiley Interdisciplinary Reviews: Climate Change*, vol. 11, no. 3, p. e642, 2020.

[13] O. Špačková and D. Straub, "Long-term adaption decisions via fully and partially observable Markov decision processes," *Sustainable and Resilient Infrastructure*, vol. 2, no. 1, pp. 37–58, 2017.

[14] R. Hui, J. Herman, J. Lund, and K. Madani, "Adaptive water infrastructure planning for nonstationary hydrology," *Advances in Water Resources*, vol. 118, pp. 83–94, 2018.

[15] S. S. Shuvo, Y. Yilmaz, A. Bush, and M. Hafen, "Modeling and simulating adaptation strategies against sea-level rise using multiagent deep reinforcement learning," *IEEE Transactions on Computational Social Systems*, vol. 9, no. 4, pp. 1185–1196, 2021.

[16] M. Pozzi, M. Memarzadeh, and K. Klima, "Hidden-model processes for adaptive management under uncertain climate change," *Journal of Infrastructure Systems*, vol. 23, no. 4, p. 04017022, 2017.

[17] K.G. Papakonstantinou and M. Shinozuka, "Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part I: Theory," *Reliability Engineering & System Safety*, vol. 130, pp. 202–213, 2014.

[18] J. R. Fischbach, D. R. Johnson, and D. G. Groves, "Flood damage reduction benefits and costs in Louisiana's 2017 Coastal Master Plan," *Environmental Research Communications*, vol. 1, no. 11, p. 111001, 2019.

[19] T. Schinko, L. Drouet, Z. Vrontisi, A. Hof, J. Hinkel, J. Mochizuki, V. Bosetti, K. Fragkiadakis, D. Van Vuuren, and D. Lincke, "Economy-wide effects of coastal flooding due to sea level rise: a multi-model simultaneous treatment of mitigation, adaptation, and residual impacts," *Environmental Research Communications*, vol. 2, no. 1, p. 015002, 2020.

[20] T. Hennequin, Y. Dong, K. Arnbjerg-Nielsen, and H. J. D. Sørup, "Life cycle assessment of a typical European single-family residence and its flood related repairs," *Journal of Cleaner Production*, vol. 228, pp. 1334–1344, 2019.

[21] NOAA, "Sea Level Trends - Battery, NY - Station ID: 8518750." `https://tidesandcurrents.noaa.gov/sltrends/sltrends_station.shtml?id=8518750`, 2025. [Accessed: April 29th, 2025].

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

[23] New York State Department of Environmental Conservation, "Appendix: Annual social cost estimates." `https://extapps.dec.ny.gov/docs/administration_pdf/vocapp23.pdf`, 2022. [Accessed: April 29th, 2025].

[24] A. Bhattacharya, G. P. Warn, K. G. Papakonstantinou, M. M. Bilec, L. McPhillips, C. E. Forest, R. Hasan, A. Sharma, and D. Chavda, "Optimal design and life-long adaptation of civil infrastructure under climate change and uncertain demands," in *ASCE Inspire 2023*, pp. 70–79, 2023.

[25] S. Rose, D. Turner, G. Blanford, J. Bistline, F. de la Chesnaye, and T. Wilson, "Understanding the social cost of carbon: A technical assessment," *EPRI technical update report (Electric Power Research Inst, Palo Alto, CA)*, 2014.

[26] W. D. Nordhaus, "Revisiting the social cost of carbon," *Proceedings of the National Academy of Sciences*, vol. 114, no. 7, pp. 1518–1523, 2017.

[27] EPA, "Standards of performance for new, reconstructed, and modified sources and emissions guidelines for existing sources: oil and natural gas sector climate review," 2022.

[28] C. P. Andriotis and K. G. Papakonstantinou, "Managing engineering systems with large state and action spaces through deep reinforcement learning," *Reliability Engineering & System Safety*, vol. 191, p. 106483, 2019.