

Leveraging Open Source Software and Parallel Computing for Model Predictive Control of Urban Drainage Systems using EPA-SWMM5

Jeffrey M. Sadler¹, Jonathan L. Goodall, Ph.D.², Madhur Behl, Ph.D³, Mohamed M. Morsy, Ph.D^{4, 5}, Teresa Culver, Ph.D⁶, and Benjamin D. Bowes¹

¹*Graduate Research Assistant, Dept. of Engineering Systems and Environment, Univ. of Virginia, 151 Engineers Way, P.O. Box 400747, Charlottesville, VA 22904, jms3fb@virginia.edu*

²*Associate Professor, Dept. of Engineering Systems and Environment, Univ. of Virginia, 151 Engineers Way, P.O. Box 400747, Charlottesville, VA 22904, goodall@virginia.edu (Corresponding Author)*

³*Assistant Professor, Computer Science Dept, madhur.behl@virginia*

⁴*Research Associate, Dept. of Engineering Systems and Environment, Univ. of Virginia, 151 Engineers Way, P.O. Box 400747, Charlottesville, VA 22904, mmm4dh@virginia.edu*

⁵*Irrigation and Hydraulics Department, Faculty of Engineering, Cairo University, P.O. Box 12211, Giza 12613, Egypt, mohamedmorsy@eng.cu.edu.eg*

⁶*Associate Professor, Dept. of Engineering Systems and Environment, Univ. of Virginia, 151 Engineers Way, P.O. Box 400747, Charlottesville, VA 22904, tbc4e@virginia.edu*

Highlights

- Open-source implementation of model predictive control for EPA-SWMM5, `swmm_mpc`
- Evolutionary algorithm used to select effective control policy at each time step
- Parallel processing of genetic algorithm significantly reduces run-time
- Control policy from `swmm_mpc` minimizes flooding and maintains target water level
- Computational cost measured for personal, high-performance, and cloud-based computers

Abstract

Active stormwater control may play an important role in mitigating urban flooding which is becoming more common with climate change and sea level rise. In this paper we describe and demonstrate `swmm_mpc`, software we developed for simulating model predictive control (MPC) for urban drainage systems using open source software (Python and the EPA Stormwater Management Model version 5 (SWMM5)). `swmm_mpc` uses an evolutionary algorithm as an optimizer and supports parallel processing. In the demonstration case, the control policy found by `swmm_mpc` for two storage units achieved its objectives of 1) practically eliminating flooding and 2) maintaining the water level at the storage units close to a target level. Although the current `swmm_mpc` workflow was feasible for a simple model using a desktop PC, a high-performance computer or cloud-based computer with more computational cores is needed for a more complex model.

Keywords

urban flooding; real-time control; parallel processing; model predictive control; stormwater; urban drainage modeling

Software Availability

The `swmm_mpc` software is open-source and available for use and improvement on GitHub at https://github.com/UVAAdMIST/swmm_mpc (note to reviewers/editor - once revisions are complete, we will make a release in the repository and include a DOI in this section). A Docker image of `swmm_mpc` is also available at https://hub.docker.com/r/jsadler2/swmm_mpc/. The demonstration model is available on HydroShare (Sadler, 2018).

1 Introduction

Researchers have predicted that storm intensity will increase on average due to climate change (Berggren et al., 2012; Neumann et al., 2015). Coastal cities have an additional challenge as sea levels rise, which makes it more difficult to drain storm runoff from streets. Coastal cities have already experienced increased flooding from high tidal events alone (Sweet and Park, 2014).

More intense storms and rising sea levels will put greater stress on urban drainage systems necessitating changes for urban drainage systems to perform at current levels. One possible adjustment is to make capital improvements such as increasing pipe size or constructing new storage units. Another option is to convert drainage systems from passive, gravity driven systems to active or “smart” systems (Kerkez et al., 2016). Active systems can increase performance of a urban drainage system at a lower cost than traditional capital improvements (Meneses et al., 2018). Actively controlling an urban drainage system does not increase the *actual* capacity of urban drainage infrastructure, but rather more efficiently uses the existing infrastructure, increasing its *effective* capacity. For example, one part of an active urban drainage system could be a valve at the outlet of a retention basin which can be automatically opened or closed based on system conditions and forecasts. With this setup, the valve could be closed more during a storm, which would utilize the available storage better than would have been possible without the valve.

For an active urban drainage system to achieve its objective (e.g., minimize flooding, reduce combined-sewer overflows), an effective management strategy is required. Management decisions for a urban drainage system include which actuators (e.g., valves and pumps) in the system should change, when to change them, and to what setting they should be changed. We refer to these decisions as a control policy (Vrabie et al., 2009; Mayne et al., 2005; Langson et al., 2004). An effective control policy for an active urban drainage system may depend on a number of factors such as antecedent moisture conditions, expected intensity and duration of oncoming rainfall, current water levels in the system, the condition of the drainage infrastructure, and other factors (e.g., tide levels in tidally influenced urban drainage systems).

A common approach for determining an effective control policy is model predictive control (MPC) (Camacho and Bordons, 2007). MPC has been used effectively in many control applications including automotive controls (Del Re et al., 2010), HVAC (heating, ventilation, and air conditioning) (Afram and Janabi-Sharifi, 2014), and other industrial applications (Qin and Badgwell, 2003). MPC has also been used effectively in urban drainage applications (Lund et al., 2018; Puig et al., 2009; Cembrano et al., 2004; Schütze et al., 2004; Gelormino and Ricker, 1994). In MPC, a process model is used to simulate the physical system and evaluate alternative control policies. Forecast data can be used as input for the simulation. During the control period, on-line optimization is performed, meaning that an optimal control policy is found and implemented at each time step (Camacho and Bordons, 2007).

Although effective for finding effective control policies, implementing MPC for a urban drainage system is non-trivial due to the dynamics within the system. The fundamental governing equations for modeling urban drainage systems are the St. Venant equations which, when considered fully, are non-linear (Tayfur et al., 1993). This makes finding an optimal control policy for urban drainage systems challenging using MPC (Darsono and Labadie, 2007). To address this dilemma, two alternative approaches are found in the literature. The first is to simplify the governing equations of the process model to a linear system. This makes the optimization problem solvable using well-established procedures such as simplex

(Nelder and Mead, 1965). Gelormino and Ricker (1994) took the approach of linearizing their system, converting their process model into a linear-time-invariant model to perform MPC for a large combined sewer system in Seattle, Washington USA.

The second approach is to retain the non-linear St. Venant equations and use a meta-heuristic to find the best control policy at each time step. In this approach, a true optimization procedure is not possible because the system remains non-linear; instead, a metaheuristic (e.g., an evolutionary algorithm (EA)) can be used (Gandomi et al., 2013). The use of a metaheuristic precludes the possibility of determining a guaranteed optimal control policy and is typically computationally expensive. The advantage of this approach, however, is that the non-linear governing equations in the process model are retained. This approach was taken by Heusch and Ostrowski (2011) who used a dynamically dimensioned search for finding the best control policy and the United States Environmental Protection Agency’s Stormwater Management Model Version 5 (EPA-SWMM5), which numerically solves the St. Venant equations, as their process model (Huber et al., 2005). Similar to Heusch and Ostrowski (2011), we have selected to follow the second approach so that the nonlinearities in the process model can remain, and to leverage EPA-SWMM5 as the process model.

EPA-SWMM5 is an attractive choice as a process model for urban drainage systems for several reasons. EPA-SWMM5 is in the public domain making it free of charge and its source code is open-source making it customizable. The model simulates a wide variety of urban drainage structures including active controls such as orifices with variable openings and pumps. EPA-SWMM5 has been used in many research applications, as well as in engineering practice to model urban drainage systems (Burger et al., 2014). Notwithstanding the wide use and utility of EPA-SWMM5 for modeling urban drainage systems, and the established utility of MPC as a successful approach for determining effective control policies, to our knowledge, there is currently no software available for performing MPC using EPA-SWMM5. Although Heusch and Ostrowski (2011) developed software that implements MPC with EPA-SWMM5, that software was closed-source and is no longer available.

This study advances the work done by Heusch and Ostrowski (2011) by creating open-source software for simulating MPC for EPA-SWMM5, *swmm_mpc*, and by demonstrating

swmm_mpc’s parallel computing capabilities. By making swmm_mpc open source, other researchers will be able to use, improve, and build from the source code. Although, the software written by Heusch and Ostrowski (2011) supported the use of parallel computing, this capability, which is critical to the usability of such software given its associated computational costs, was never demonstrated or tested in the literature.

swmm_mpc was written in the Python programming language. Several third-party Python packages were necessary for the success of this project including pyswmm (<https://github.com/OpenWaterAnalytics/pyswmm>) and the Distributed Evolutionary Algorithms for Python (DEAP) (<https://github.com/DEAP/deap>). To evaluate swmm_mpc, it was applied to a use case model with two active control devices. The swmm_mpc results were compared to the results from a rules-based approach and a scenario with no active control. The swmm_mpc software was run on a desktop personal computer (PC), a high-performance computer (HPC), and a rented, cloud-based machine to demonstrate and test the parallel processing capability of the software.

The remainder of this paper describes the methods used to implement swmm_mpc including a description of the MPC workflow and the interaction and role of the third-party Python libraries. The use case model is then described and the results of the evaluation are presented and discussed. As part of the results and discussion, the benefits of parallelization and the use of a high-performance and cloud-based computing for running swmm_mpc are quantified and discussed.

2 Methods

2.1 Overview of MPC for urban drainage systems

MPC for a urban drainage system consists of three main components as shown in Figure 1. The first component is the physical system, including the system states and system controls. The system states include hydraulic states such as water levels at system nodes and flow rates in system pipes, and hydrologic states such as watershed soil moisture and runoff. In a real system, these states would come from real-time sensors. The system controls are actu-

ators that accept and implement the settings resulting from the MPC process at each time step.

The second component in MPC is a process model used to simulate the future states of the urban drainage system. The process model uses the states read from the urban drainage system as its initial states. The process model also takes future model inputs such as rainfall or tide level. Given the current state of the system and future disturbances, the process model is used to evaluate the effectiveness of control policy candidates.

A control policy consists of one setting for each actuator, for each control time step, for the duration of the control horizon. An individual setting can be a number as would be the case for a valve where the number would correspond to the percent open of the valve. An individual setting can also be a binary setting as would be the case for a pump that can either be on or off. As an example, consider a system with a variably opening valve and an on-off pump with a control horizon of one hour and a control time step of 15 minutes. A control policy for this system would consist of two arrays, an array of numbers between 0 and 1 to specify the percent open of the valve should be, and an array of “on” or “off” to specify the setting of the pump. Both arrays would have four settings (one for each time step in the control horizon of one hour).

To evaluate the effectiveness of a given control policy, the settings in the policy are applied to simulated actuators in the process model and the process model simulation is executed. At the end of the simulation, a cost is determined for the policy. The cost is based on a user-defined cost function. In this study, we consider mainly the cost resulting from flooding but other costs could be considered within this general framework including the costs of combined sewer overflows (CSO) and water quality. The cost may also be a factor of other process model outputs such as deviation from target water levels at certain points (Schütze et al., 2004).

The third component of MPC for a urban drainage system is an optimization routine to determine the best control policy for the system. Using the process model to assign a cost to a given control policy, the optimization procedure seeks to find the control policy that incurs the smallest cost. If the process model is linear, a true optimum can be found using tradi-

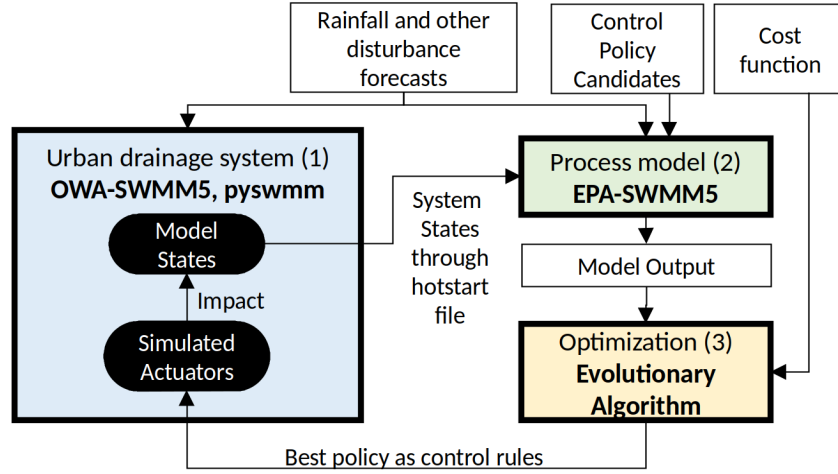


Figure 1: Main components of MPC in swmm_mpc system

tional optimization procedures like simplex (Nelder and Mead, 1965). If the process model is non-linear, other approaches must be taken such as using a metaheuristic to find an effective control policy (Gandomi et al., 2013).

In summary, the chronological workflow for MPC for a urban drainage system is: 1) system states are read from the physical system, 2) using the system states as initial conditions and future disturbances as input, a process model is used to evaluate control policies, 3) the best control policy is selected through an optimization procedure, and 4) the best control policy is implemented in the real system. Although the best control policy is obtained for the entire control horizon, only the first step in the control policy is used since the procedure re-optimizes at every control time step.

2.2 MPC for SWMM5: swmm_mpc

Implementation of the parts of MPC using Python and SWMM5 was done in the swmm_mpc Python package. The software simulates online MPC for an urban drainage system using SWMM5 as the process model and as the simulated physical system. The current system could also be used in an offline mode to find a control policy for a forecast storm event beforehand.

2.2.1 Simulated urban drainage system: OWA-SWMM5 and pyswmm

In `swmm_mpc` an enhanced version of SWMM5, OWA-SWMM5 (<https://github.com/OpenWaterAnalytics/Stormwater-Management-Model>), was used via an accompanying Python library, `pyswmm` (<https://github.com/OpenWaterAnalytics/pyswmm>) to simulate the physical urban drainage system. Both OWA-SWMM5 and `pyswmm` were developed and are distributed by Open Water Analytics.

OWA-SWMM5 and `pyswmm` provide three key functionalities needed to simulate the online optimization procedure required by MPC. First, unlike when a simulation is run via EPA-SWMM5, when using `pyswmm`, custom Python routines can be executed between each time step of the simulation. This is critical to `swmm_mpc` because at each time step in the MPC simulation workflow three processes occur: 1) the states from the simulated urban drainage system are read and transferred to the process model; 2) the metaheuristic is run; and 3) the best policy found by the metaheuristic needs is implemented in the simulated urban drainage system. Using `pyswmm`, Python code can be run to perform each of these processes at each control time step.

Second, `pyswmm` enables the transfer of system states at each time step from the simulated urban drainage system to the process model. This is accomplished through a hotstart file. A SWMM5 hotstart file contains all of the hydraulic and hydrologic states of the model at the time in the simulation when the hotstart file is saved. When a hotstart file is read into a simulation, that simulation's initial hydraulic and hydrologic states are the states represented in the hotstart file. This functionality is well-suited to transfer the states of the simulated urban drainage system to the process model in the `swmm_mpc` workflow.

Using EPA-SWMM5, a hotstart file can be saved only at the end of a simulation. This is a critical limitation because in MPC the system states need to be transferred at every time step. To address this limitation, we added new functionality to OWA-SWMM5 and `pyswmm` to enable hotstart files to be saved at any point in a SWMM5 simulation executed using `pyswmm`. This functionality allowed the system states of the simulated urban drainage system to be transferred to the process model at each time step.

Third, through pyswmm the best control policy found by the metaheuristic can be implemented at each time step. This is done using pyswmm to change the settings of the actuators in the model during the simulation. When a simulation is initialized in pyswmm, each object in a SWMM5 model (every node, link, subcatchment, etc.) can be read into a Python object via its element ID as defined in the SWMM5 input file. Each of these Python objects has attributes that can be read (e.g., depth at a node and flow in a link). Actuators in the model read into Python objects also have the “target_setting” attribute that can be written. To implement a control setting for an actuator via pyswmm, its “target_setting” is set to the first setting in the best control policy.

2.2.2 Process model: EPA-SWMM5

In addition to representing a real urban drainage system, SWMM5 was used as the process model. However, in contrast to using OWA-SWMM5 to simulate the physical urban drainage system, the standard EPA-SWMM5 was used as process models. This was necessary because the current version of pyswmm cannot run more than one simulation at a time. This is a functionality needed in swmm_mpc because at each time step during the simulation of the urban drainage system, at least one process model simulation is run in a predictive fashion to evaluate control policy candidates. EPA-SWMM5, unlike pyswmm, supports multiple simulations being executed simultaneously.

2.2.3 Active controls in EPA-SWMM5

EPA-SWMM5 simulates the active control of certain hydraulic structures including pumps, orifices, and gates. Each of these structures has a setting that can be assigned. For example, the setting for an orifice is a decimal number between 0 and 1 which corresponds to the percent open of the orifice (e.g., a 0.5 setting would mean the orifice was 50%). The user can also define an amount of time for a structure to implement a change in setting. This “time to change” parameter in EPA-SWMM5 represents the delay seen in reality for changing an actuator’s setting.

Changing controls during an EPA-SWMM5 simulation is done using one or more control

```

Rule 01
If NODE J1 DEPTH <2
THEN ORIFICE 01 SETTING = 0.6

```

Figure 2: Example of a control rule in SWMM5

rules (see example in Figure 2). A control rule is specified in the SWMM5 input file before the simulation begins and consists of four parts. The first two parts of a control rule are the rule name and the condition. In the example, the rule name is “R1”. The condition is “IF NODE J1 DEPTH <2”, meaning that the program will check if the depth at the node with the ID of “J1” is less than 2 (the units being defined globally in the model input file as feet or meters). In EPA-SWMM5 the condition can be the state at any link or node and can also be related to global simulation states such as the model simulation time. The third part of a control rule defines which structure(s) should change if the specified condition is met. In the example the structure that will change is “ORIFICE R1”. Finally, the fourth part of the rule defines the setting to which the structure should change. In the example, this is “0.6”, meaning that if the condition is met, the orifice should be set to 60% open.

In `swmm_mpc` a control policy is a time series of control settings (one control setting per control time step for the control duration). This is implemented in EPA-SWMM5 as a set of control rules. Since a control policy in MPC is a time series, each control rule’s condition is based solely on the model’s simulation time in decimal hours. For example, Figure 3 shows a control policy of four settings (0.2, 0.4, 0.5, and 0.2) for “ORIFICE R1” at a 15-minute control time step implemented as control rules. This text would be written to the EPA-SWMM5 process model input file under the “CONTROLS” heading.

2.2.4 Metaheuristic: evolutionary algorithm

Because we used EPA-SWMM5 as a black-box process model, a metaheuristic was used in place of a true optimization procedure to find an effective control policy at each time step in the MPC run. We chose an evolutionary algorithm (EA) for the metaheuristic since it has been shown to be successful in other urban drainage control applications (Zimmer et al., 2015, 2018) and it’s inherent propensity for parallelization Maier et al. (2014). An EA be-

```

Rule R1
If SIMULATION TIME < 0.25
THEN ORIFICE O1 SETTING = 0.2

Rule R2
If SIMULATION TIME < 0.50
THEN ORIFICE O1 SETTING = 0.4

Rule R3
If SIMULATION TIME < 0.75
THEN ORIFICE O1 SETTING = 0.5

Rule R4
If SIMULATION TIME < 1.0
THEN ORIFICE O1 SETTING = 0.2

```

Figure 3: Example implementation of control policy as set of control rules

gins with an initial population of individuals where, in our case, each individual is a control policy. A fitness score (or conversely a cost) is assigned to each individual in the population and certain individuals are selected to survive into the next generation based on their fitness score. Mechanisms for improving the fitness of the individuals from one generation to the next mimic natural processes including cross-over and mutation (Maier et al., 2014). The process of selection and improvement is repeated from generation to generation until a stopping criteria is met. Common stopping criteria include a user-defined number of generations or an acceptably low rate of improvement from one generation to the next. The use of an EA requires several user-defined parameters including the number of individuals in the initial population, the cross-over rate, the mutation rate, and the stopping criteria.

In `swmm_mpc`, policies are transcribed to a string of bits in the genetic algorithm. The settings for orifices and weirs are represented by three bits with eight total possibilities (0/8 open, 1/8 open, ..., 8/8 open). We limited the settings of orifices and weirs to three bits to limit the solution space explored by the algorithm. The setting of a pump is represented by only one binary digit (1 or “ON”, 0 or “OFF”).

Since the EA searches for the policy that incurs the minimum cost, the way in which a cost is assigned to each individual control policy is very influential on the EA’s effectiveness. In `swmm_mpc`, the cost of a control policy is determined using the process model and a cost function. First, each individual control policy is implemented in the process model input file as a set of control rules as described above. Once the control policy is implemented, the

EPA-SWMM5 model is executed. Elements of the model output resulting from the process model execution become input for the cost function. The cost function used in `swmm_mpc` is

$$Cost = \alpha(\mathbf{a} \cdot \mathbf{v}(\mathbf{u}, \mathbf{x})) + \beta(\mathbf{b} \cdot \mathbf{d}(\mathbf{u}, \mathbf{x})) \quad (1)$$

where \mathbf{a} , \mathbf{v} , \mathbf{b} , \mathbf{d} are each 1-dimensional vectors, \mathbf{u} and \mathbf{x} are 2-dimensional vectors, and α and β are scalars. The members of \mathbf{a} are user-defined weight values for flooding at any node in the system and the members of \mathbf{v} are the flood volumes at each node over the entire simulation as calculated by the process model. The members of \mathbf{b} are user-defined weights for deviation from user-defined target water levels at each node in the system and the members of \mathbf{d} are the average absolute deviations from target water levels again as calculated by the process model over the entire simulation. \mathbf{u} represents the control policies for all controls, for each time step. \mathbf{x} represents the system states. α , and β , are user-defined constants used to scale and give overall weights to flooding costs compared to deviation costs. Typically weights for the components of the cost (or objective) function sum to 1 and can include a scaling factor to account for variables in different units or scales (Kim and de Weck, 2005). In this formulation, α and β include both the weight and the scaling factors for the objectives.

We intentionally made this cost function flexible so that users can customize it to meet their objectives which may vary between use cases. A cost for flooding is obviously important as that is a major concern for many communities and the prevention of which is one of the main purposes for urban drainage systems. We also included a cost from deviations for target water levels because, in certain cases, it is desirable to maintain water levels close to a target depth. For example, it may be important to keep a certain amount of water in a retention pond for aesthetic and/or ecological purposes. Although the cost function is flexible, when implemented in `swmm_mpc`, the user need only define what is important to the specific application. For example, default for \mathbf{a} is a vector of all 1's. When one node is specified, the weight of any unspecified node becomes zero. The default for \mathbf{b} is all zeros, since the user has to specify a target depth for a given node.

To execute EAs we used the Distributed Evolutionary Algorithms for Python (DEAP)(<https://github.com/DEAP/deap>) library. An advantage of EAs is that they can easily be run in parallel since they perform many independent evaluations (Maier et al., 2014). In DEAP parallel processing is supported through integration with the built-in “multiprocessing” Python library.

2.2.5 swmm_mpc workflow

The MPC workflow in `swmm_mpc` was implemented using three main Python functions (see Figure 4). The function in the workflow called by the user is “`run_swmm_mpc`.” This function runs the MPC workflow and calls the two other main functions. “`run_swmm_mpc`” takes 13 user inputs as shown in Table 1. Through these inputs, the user specifies the model input file to that represents the urban drainage system, the control inputs (i.e., which controls to find a policy for, the control time step, and the control horizon), and EA parameters (e.g., number of generations, cost function parameters).

The most complex of the user-supplied arguments are “`target_depth_dict`” and “`node_flood_wgt_dict`” (see Snippet 1 for examples). These two arguments define the \mathbf{a} and \mathbf{b} variables in the cost function. Additionally, the “`target_depth_dict`” argument is used to determine \mathbf{d} . These arguments map from Python data structures to the mathematical variables in the cost function. The “`target_depth_dict`” argument is a dictionary whose keys are node ids and whose values are dictionaries. The inner dictionary has two keys, the target depth of the node and the weight of the cost for deviations from the weight at the node. In Snippet 1, the “`target_depth_dict`” specifies that the target depths of Nodes St1 and St2 are 4.0 and 3.5, respectively. The weights are also specified: deviation from the target depth at Node St1 will be twice as costly as deviation from Node St2. The “`node_flood_wgt_dict`” is a simpler dictionary, the keys of which are node ids and the values are weights. In Snippet 1, the “`node_flood_weight_dict`” specifies that flooding at Node J1 is five times costlier than flooding at Node St1. Note that if one or more node is included in the “`target_depth_dict`” or the “`node_flood_wgt_dict`”, other nodes are not included in the cost calculation (in terms of the cost function, the corresponding weights in \mathbf{a} and \mathbf{b} are zero). This is shown in Snippet 1,

the weight of deviations from a water level at Node J1 and the weight of flooding at Node St2 would both be zero since they are not included in the dictionaries.

Snippet 1: Examples of “target_depth_dict” and “node_flood_wgt_dict”

```
target_depth_dict = {"Node St1": {"target": 4.0, "weight": 2},
                    "Node St2": {"target": 3.5, "weight": 1}}

node_flood_weight_dict = {"Node J1": 1, "Node St1": 0.2}
```

In the “run_swmm_mpc” function, the SWMM5 model simulating the urban drainage system is run step by step via pyswmm. At the beginning of the simulation, the SWMM5 input file representing the urban drainage system is copied. This copy serves as the input file used for the process model. To ensure that the states and simulation periods of process model remain in sync with the simulated urban drainage system, at each time step a hotstart file from the urban drainage system simulation is saved and then used as the initial states for the process model. The process model’s simulation start date and time are also updated to match the urban drainage system simulation’s current date and time.

Once the process model’s simulation start date and time is same as the simulated urban drainage system and the hotstart file of the simulated urban drainage system is saved, the “run_ea” function is called. The “run_ea” function initiates the EA which starts by creating an initial population of individual control policies. In our case, an individual is a 1-dimensional vector, each member of which is a setting for an individual actuator for one control time step.

The initial population for the first time step is a group of random individuals. For subsequent time steps, elitism is used where the best policy found in the previous time step is used to seed the initial population of the current time step. For example, consider a control policy for i control time steps for j controls represented by a ixj vector, \mathbf{u} . For the first time step, the settings for time 0, \mathbf{u}_0 , through time i , \mathbf{u}_i , are found, however only \mathbf{u}_0 is implemented. Because settings that minimize cost for times $1, 2, \dots, i$ ($\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_i$) have also been found by the optimizer, in swmm_mpc, these settings are used as the basis for seeding the

population for the next time step. Because the seed population is for the following time step, \mathbf{u}_1 serves as \mathbf{u}_0 , \mathbf{u}_2 serves as \mathbf{u}_1 , etc. and \mathbf{u}_i is set randomly. This base individual is then varied randomly to create several similar individuals. Some randomly created individuals are also added to the population for the next generation to create variety in the policies and to ensure the number of policies is equal to the initial population size specified by the user.

The control policies initiated in the “run_ea” function are input into the third main function, “evaluate”. The evaluate function makes a copy of the process model input file and the input hotstart file. To avoid file naming conflicts, a random string is appended to the hotstart and input file names. The control policy is then implemented in the newly created input file by adding corresponding control rules. Once the control policy is implemented in the input file, the simulation is executed with EPA-SWMM5. When the simulation run is completed, the “evaluate” function parses the output file to determine \mathbf{v} and \mathbf{d} in the cost function. The policy’s cost can then be determined since the remaining cost function parameters (α , \mathbf{a} , β , and \mathbf{b}) are user-defined. The evaluation of an individual control policy is independent of all others, therefore, the “evaluate” function is the part of the workflow that is parallelized through Python’s “multiprocessing” module.

Using the cost that has been assigned to each policy, the “run_ea” function selects the best individuals to retain in the population of policies for the next generation. After the user-defined number of generations are complete, the best policy found by the EA is implemented in the simulated urban drainage system in the “run_swmm_mpc” function. The policy is also saved and used to seed the population of the next time step. Finally, the setting for that time step is recorded so that at the end of the simulation, the best control policy for the entire simulation time is saved. If ever the cost at a time step is zero, the algorithm stops since the performance cannot improve further.

2.3 System demonstration

To demonstrate the utility and functionality of swmm_mpc, three cases (A, B, and C) with increasing complexity were implemented in a simple SWMM5 model. Three control scenarios (passive, rules-based, and MPC) were applied to each of the cases. The computational costs

Table 1: User inputs for “run_swmm_mpc” function

Parameter name	Data type	Description	Default value
inp_file_path	String	File path to SWMM5 input file (.inp)	N/A
control_horizon	Number	Control horizon in hours	N/A
control_time_step	Number	Control time step in seconds	N/A
control_str_ids	List of strings	IDs of control structures to be adjusted	N/A
work_dir	String	Path to directory where temporary files will be stored	N/A
results_dir	String	Path to directory where results should be stored	N/A
target_depth_dict	Dictionary	IDs of nodes and corresponding target depths and relative penalty weights	Null
node_flood_weight_dict	Dictionary	IDs of nodes and corresponding relative penalty weights	Null
flood_weight	Number	Overall weight of flood penalties	1
dev_weight	Number	Overall weight of deviation penalties	1
ngen	Number	Number of generations that the GA should perform	7
nindividuals	Number	Number of individuals in the initial GA population	100
run_suffix	String	suffix to be appended to results file	

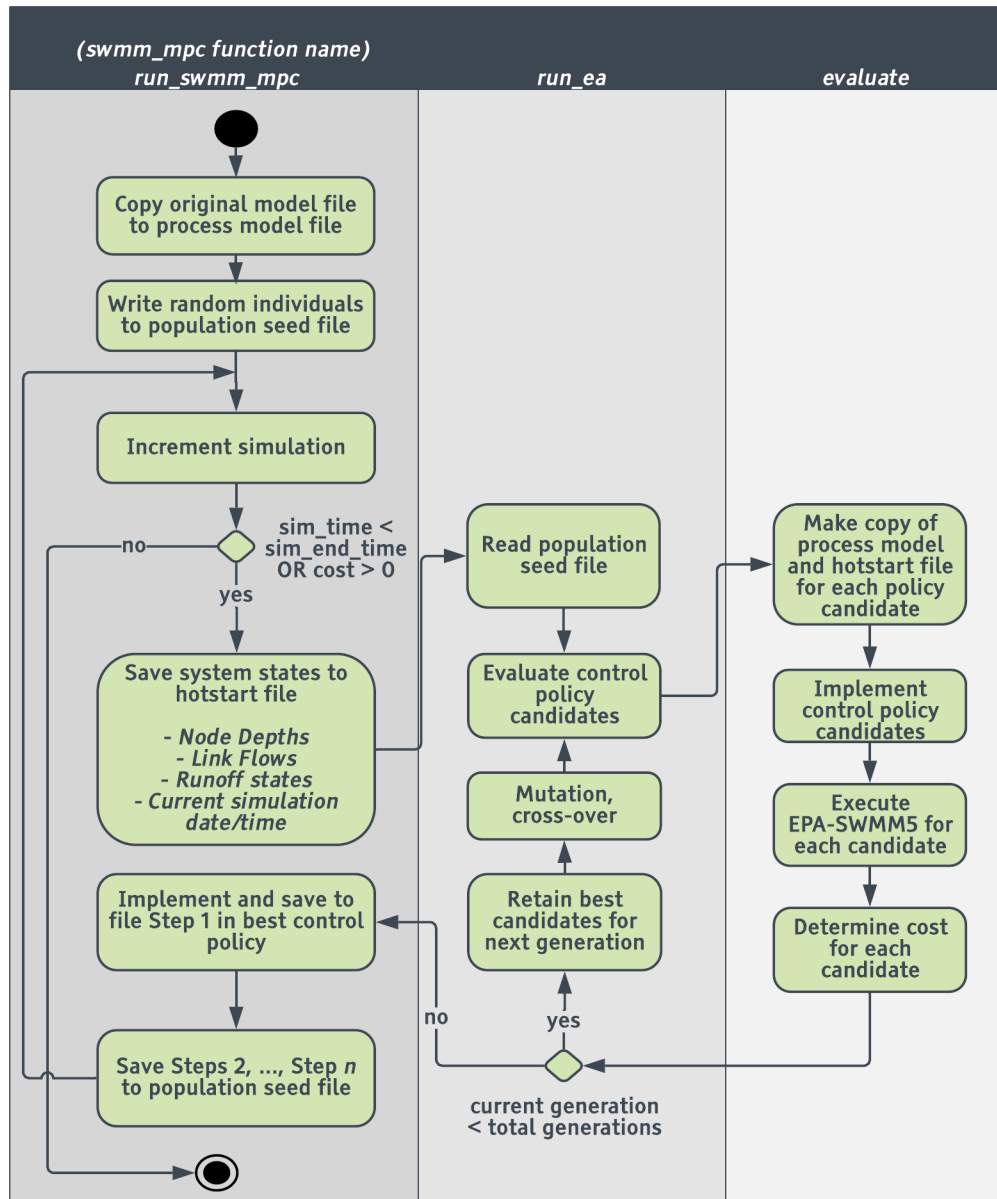


Figure 4: Activity diagram of MPC implementation using Python

of running `swmm_mpc` were also quantified.

2.3.1 Demonstration model and rainfall event

A schema of the model used to demonstrate `swmm_mpc` is shown in Figure 5 and the model properties are given in Table 2. Since the intent of this paper is to demonstrate the software, a simple model was used. However, any SWMM5 model with controllable features can be used with `swmm_mpc`.

The demonstration model has two subcatchments, S1 and S2. For the example use case, we used a 1-year, 12-hour design storm for Norfolk, Virginia, a coastal city that experiences frequent flooding (Mitchell et al., 2013). The design storm (see Figure 6) had 78.2 mm of total rainfall (Bonnin et al., 2018) with an SCS Type II temporal distribution (Mockus, 2012). The use of design storms is typical when designing and modeling stormwater infrastructure (for example Villarreal et al. (2004) and Hatt et al. (2009)). To introduce some spatial variation in the model, the rain for S2 started 18 minutes after the start of the rain event for S1. S2 is smaller but flashier with a smaller width and a larger percent of impervious surfaces. This is manifested in a higher peak flow shown in Figure 6.

The subcatchment S1 and S2 drain directly to two storage units in parallel, St1 and St2, respectively. In SWMM5 models, storage units are generic and are used to represent both natural storage facilities such as a pond as well as man-made facilities such as an underground tank or retention pond. Two orifices (R1 and R2) control the flow out of two storage units. The two orifices from the storage units meet and flow through a junction, J1, before leaving the system through the outfall. The model simulation time was 24 hours and the routing time step was ten seconds in the MPC and passive scenarios and five seconds in the rules-based scenario.

2.3.2 Demonstration cases

We implemented and compared the three control scenarios for three use cases with different objectives and model constraints increasing in complexity (see Table 3). In Case A, the simplest of the three cases, the only objective was to minimize flooding. In Case B, an addi-

Table 2: Properties of demonstration model

Subcatchment properties				
Name	Area (hectares)	Width (km)	Slope (%)	Impervious (%)
S1	32	0.46	0.5	30
S2	20	0.61	0.5	45
Storage unit properties				
Name	Surface Area (m ²) (constant)	Initial Depth (m)	Max Depth (m)	Bottom Elev (m)
St1	6039	0.5	1.5	0.91
St2	4645	0.5	1.5	0.91
Pipe properties				
Name	Length (m)	Diameter (m)	Roughness	
C1, C2	122	0.3	0.01	
Node properties				
Name	Max depth (m)	Bottom Elev (m)		
J2	1.5	0.91		
J1	0.6	0.30		
Outfall	NA	0		
Orifice properties				
Name	Height (m)	Discharge Coef- ficient		
R1 & R2	0.61	0.65		

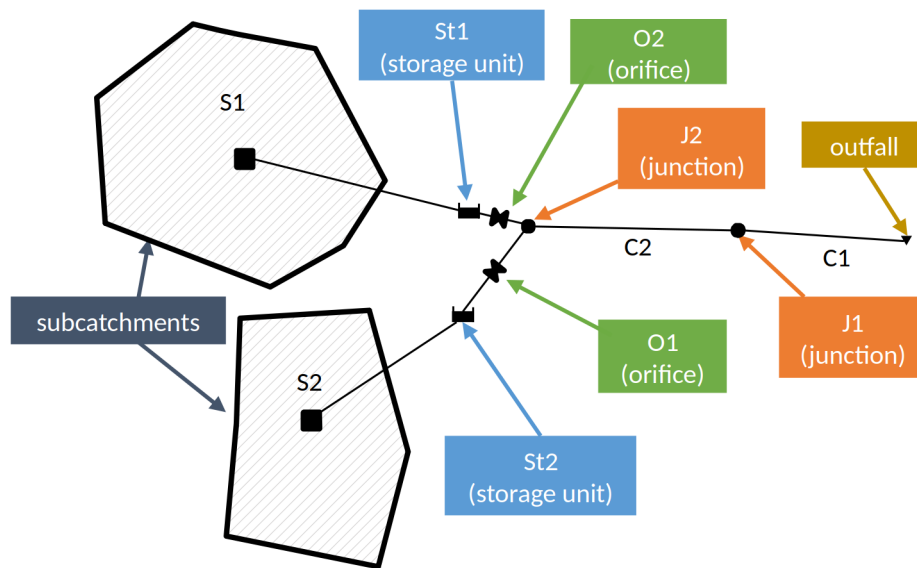


Figure 5: Demonstration model schema

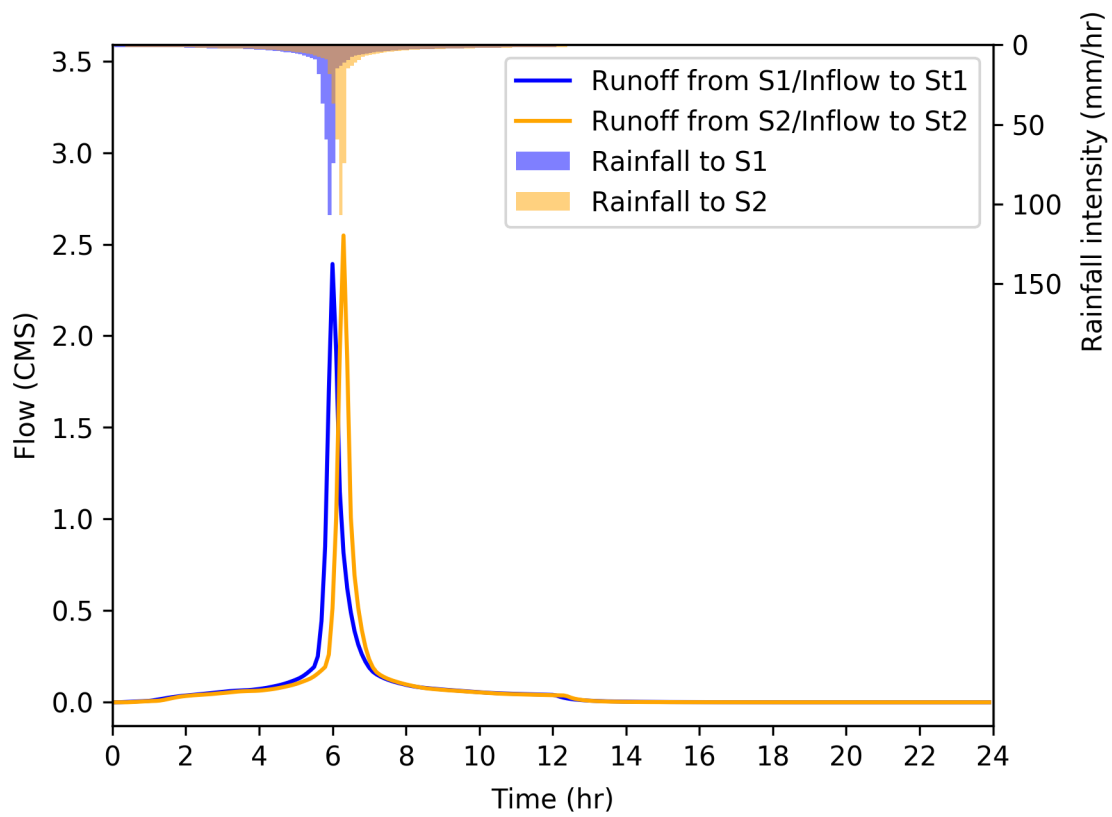


Figure 6: Design storm input to use case study

Table 3: Differences in MPC cost function parameters between test cases

Cost function parameter (description)	Case A	Cases B & C
α (overall flood weight)	1	1000
\mathbf{a} (individual node flood weight [St1, St2, J1, J2])	[1, 1, 1, 1]	[1, 1, 1, 1]
β (overall deviation weight)	0	0.5
\mathbf{b} (individual deviation weight [St1, St2, J1, J2])	NA	[1, 1, 0, 0]
(target depths (m) [St1, St2, J1, J2])	NA	[0.5, 0.5, NA, NA]

tional objective of maintaining a target water depth of 0.5 m at the storage units was added. In Case C, the objective was the same as Case B, but a tidal boundary condition at the out-fall was added (see Figure 7).

Since Cases B and C are multi-objective, α and β had to be selected based on the relative importance and the relative scale of each. The units of \mathbf{v} were 10^6 gallons in the SWMM5 demonstration model, much smaller in magnitude in our case than the deviations from target water levels which were in units of feet. To scale the two variables and emphasize minimizing flooding at Node J1 over minimizing deviations from the target water depths at the storage units, we set the cost of flooding, α , 2000 times larger than the cost of deviations, β (1000 compared to 0.5) for Cases B and C.

2.3.3 Control scenarios

Scenario 1: Passive In this scenario there is no active control. The outlets are 100% open at all times.

Scenario 2: Rule-based control For this scenario, simple logical rules controlled the orifice openings and therefore the discharge from the storage units. In practice, such rules can be based on experience and knowledge gained by local stormwater personnel over time. Although, heuristic-based rules alter a dynamic, actuated system, the rules themselves are static, meaning that they do not change for the duration of an event. Furthermore, the rules do not adjust based on forecast conditions.

There were two sets of rules used in the demonstration cases. For Case A, where the only objective was to reduce flooding, the rules (see Figure 8) specified to open the orifice com-

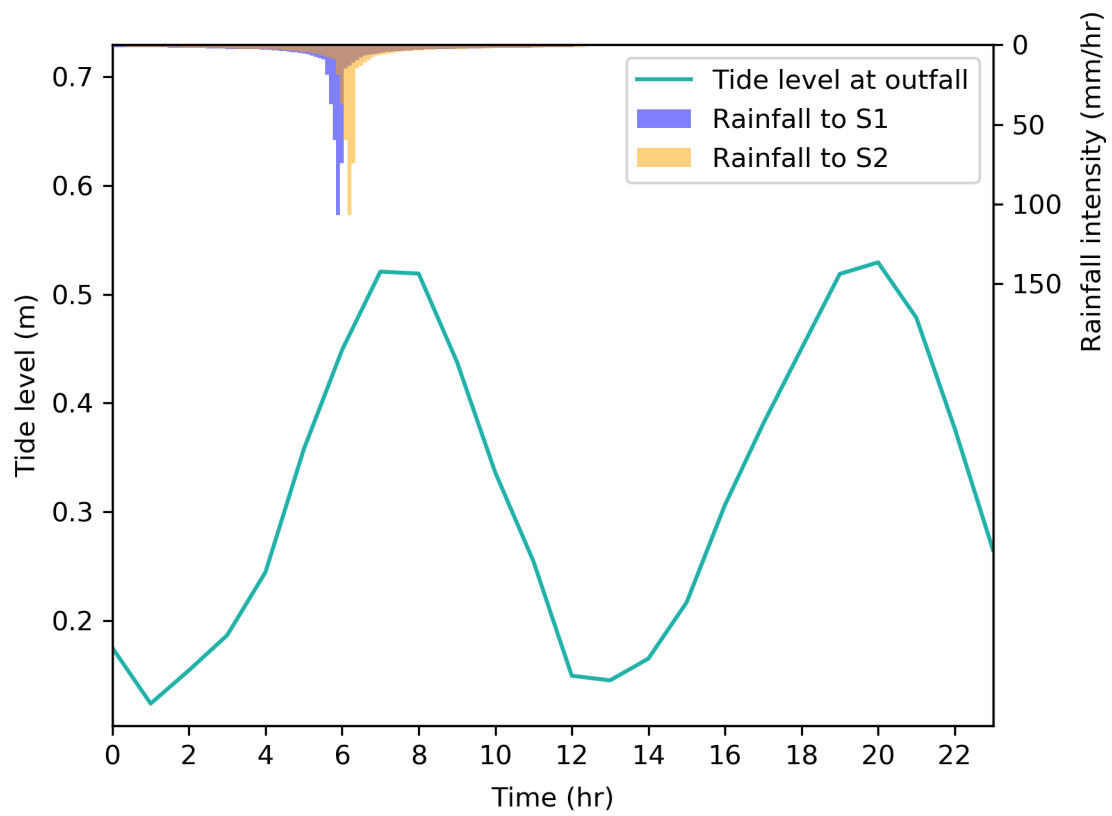


Figure 7: Tidal boundary condition at outfall used in Case C

```

Rule R1
If NODE J1 DEPTH < 0.49
THEN ORIFICE R1 SETTING = 1
AND ORIFICE R2 SETTING = 1

Rule R2
If NODE J1 DEPTH > 0.49
THEN ORIFICE R1 SETTING = 0
AND ORIFICE R2 SETTING = 0

```

Figure 8: Control rules for Case A

pletely as long as the depth in the downstream node, J1 was below 0.49 m (80% of the maximum depth). Additionally, to avoid overtopping, if the storage units reached 1.49 m (98% of the maximum depth) the orifices were to open to avoid overtopping of the storage units.

The rules for Cases B and C were the same. These rules (see Figure 9) were similar to the Case A rules, however, in addition to reducing flooding, the rules for Cases B and C also took into account the target depth of 0.52 m. In these rules the orifices were open when the depth at Node J1 was below 0.49 m and when the depth in the storage units was above the target water depth. The orifices were closed when the depth at Node J1 was above 0.49 m or when the depth in the storage units was below the target depth. The control rules scenarios for each case were run via pyswmm with a control time step of 5 seconds.

Scenario 3: MPC The MPC control policy was found using `swmm_mpc` as described in Section 2.2.5 above. One advantage of MPC over the rule-based control is the ability to adjust the actuators based on forecast conditions. For the use cases, we used a control time step of 15 minutes and a control horizon of one hour. Therefore, with two controls, a single control policy consisted of a vector of eight values (2 controls x 4 control steps per hour x 1 hour) (see Table 4). The number of individuals for the initial population of policies in the genetic algorithm was 120 and the number of generations was 8.

2.3.4 Use of parallel, high-performance, and cloud computing

The EA used for selecting the best control policy is computationally expensive and therefore, some analysis of computational costs for executing the `swmm_mpc` workflow was performed.

```

Rule R1
If NODE J1 DEPTH < 0.49
AND NODE St1 DEPTH > 0.52
THEN ORIFICE R1 SETTING = 1

Rule R2
If NODE J1 DEPTH < 0.49
AND NODE St2 DEPTH > 0.52
THEN ORIFICE R2 SETTING = 1

Rule R3
If NODE J1 DEPTH > 0.49
OR NODE St1 DEPTH < 0.52
THEN ORIFICE R1 SETTING = 0

Rule R4
If NODE J1 DEPTH > 0.49
OR NODE St2 DEPTH < 0.52
THEN ORIFICE R2 SETTING = 0

Rule R5
If NODE St1 DEPTH > 1.49
THEN ORIFICE R1 SETTING = 1

Rule R6
If NODE St2 DEPTH > 1.49
THEN ORIFICE R2 SETTING = 1

```

Figure 9: Control rules for Case B and C

Table 4: Control parameters for demonstration cases

Parameter	Value
Control Horizon (hr)	1
Control Time Step (hr)	0.25
Num controls	2

Table 5: Specifications of computational resources used for demonstration model

	PC	HPC	GCP
Max. # cores	8	28	64 (tested up to 32)
CPU speed	3.60 GHz	2.4 GHz	2.0 GHz
Processor type	Intel i7	Intel Xeon	Intel Xeon
RAM	16 GB	128 GB	7-120 GB (depending on # of CPUs)

The wall-clock times for Case A (the more complex of the two cases) were compared when using a typical personal computer (PC) and the University of Virginia’s high-performance computing (HPC) system, Rivanna. Recognizing that many (likely most) municipalities will not have HPC resources available to them, we also explored the use of a commercial cloud computing service for running `swmm_mpc`. These services, such as Amazon Web Services, Google Cloud Platform, or Microsoft Azure, allow users to rent large, powerful computers, charging only for the time that the computers are being used. To explore the option of renting a cloud-based machine, we also executed Case A through Google Cloud Platform (GCP). The number of cores available, RAM, and processor speeds of the PC, HPC, and GCP machines are listed in Table 5. Case A was run with varying number of cores on each platform.

3 Results and Discussion

3.1 Results from demonstration cases

Figures 10, 11, 12 show the results of the three control scenarios applied to the three demonstration cases. In all cases, Cases A, B, and C, there are times where flooding occurred at one or more nodes. Flooding occurs in SWMM5 when the depth of water at a given node exceeds the maximum depth of that node. When this occurs, the depth in SWMM5 at that node is recorded as the maximum depth which is why some of the depth values in the Figures 10, 11, 12 appear capped.

Figure 10 shows the results from the three control scenarios for Case A. In the rule-based and MPC scenarios, the control policies kept the valves closed more on average, thus retaining more water in the storage units and preventing flooding at Node J1 (see Figure 10

A and B). The water level at St2 reaches much higher values in the rules-based and swmm_mpc scenarios (peak of 1.16 m (3.80 ft) and 1.14 m (3.73 ft)) compared to the passive (peak of 1.05 m (3.45 ft)). In Case A, the rules-based control and the swmm_mpc scenario eliminated flooding altogether. This was the simplest control case and the algorithm in swmm_mpc was able to find a policy that eliminated flooding by hour 3. This policy resulted in a cost function equaling zero which stopped the algorithm from running. Therefore, after hour 4, the orifice settings were unchanged in the swmm_mpc scenario. The rules-based scenario in Figure 10 causes oscillation around the 80% full depth at J1 (0.49) as a result of the control rules.

For Case B, a target water depth at St1 and St2 of 0.52 m was introduced and the initial depth at St1 and St2 was 0.52 m. With more water in the system to begin with, much more flooding occurred in general compared to Case A and neither the rules-based nor the swmm_mpc control was not able eliminate flooding. The swmm_mpc scenario, however, reduced flooding by 90% compared to the passive where the rules-based reduced flooding by only 54%. The main reason for the better performance by the swmm_mpc was its ability to take action proactively and prioritize reducing flooding over maintaining the target depth. The swmm_mpc algorithm allowed the water level at St2 to go well below the target depth at the beginning of the simulation to create more storage for when the peak of the runoff arrived. Additionally the swmm_mpc policy held more water back in St1 than the rules-based scenario, more fully utilizing the storage available in the larger of the two storage units. Because the rules in the rules-based scenario maintained the water at the target depth in St2, there was not enough storage available to handle the peak of the runoff and the storage unit was overtopped.

The addition of the tidal boundary condition in Case C caused increased flooding for all three control scenarios. Again, the swmm_mpc was able to significantly reduce flooding, resulting in a 74% reduction compared to the passive scenario. The swmm_mpc policy held even more water back, and held it back for longer in St1 in Case C compared to Case B. This reduced flooding in J1 during the second high tide which occurred around 19 hours. The swmm_mpc policy did not outperform the rules-based scenario by as much as Case B

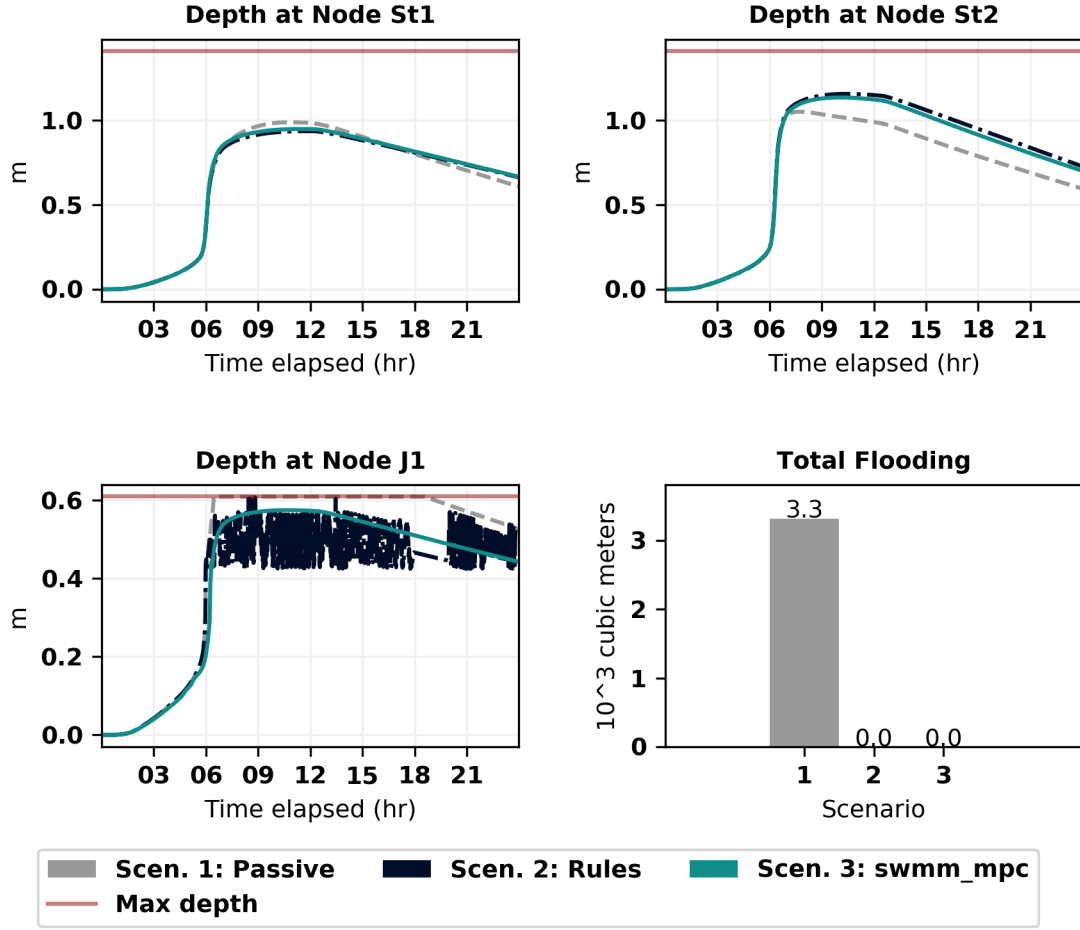


Figure 10: Depth and flooding at system nodes for Case A

- the rules-based scenario reduced flooding by 58% in Case C. A possible explanation for is that, with the additional water volume from the tidal condition, the system may be reaching its physical limits in terms of total available storage.

3.2 Computational cost of swmm_mpc

Figure 13 shows the wall-clock times for executing swmm_mpc for Case C on a PC, an HPC, and GCP machines with a varying number of processing cores. The simulation had 96 control time steps (15 minute resolution for 24 hours). If used for online MPC in a real case, the wall-clock time required for one time step would need to be less than the time step itself, otherwise, the setting for the next time step would not be determined before it would need

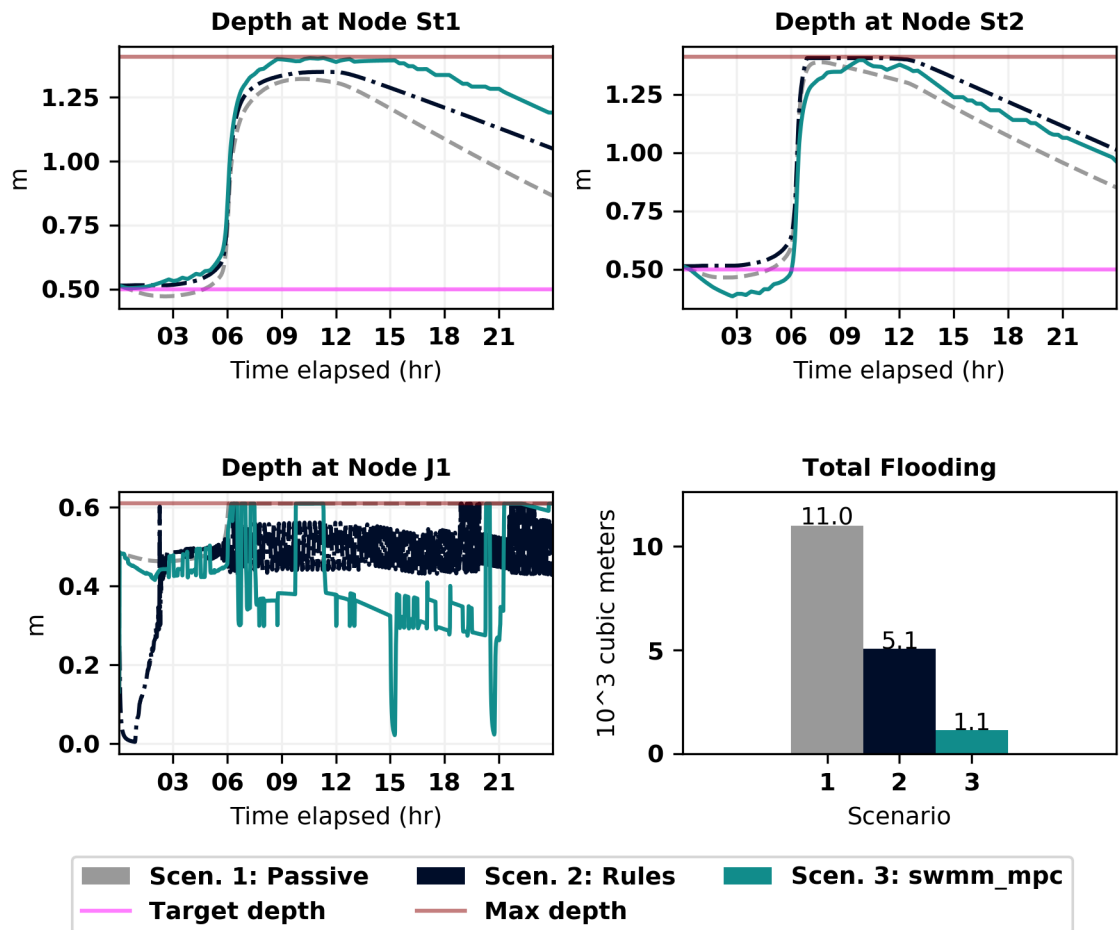


Figure 11: Depth and flooding at system nodes for Case B

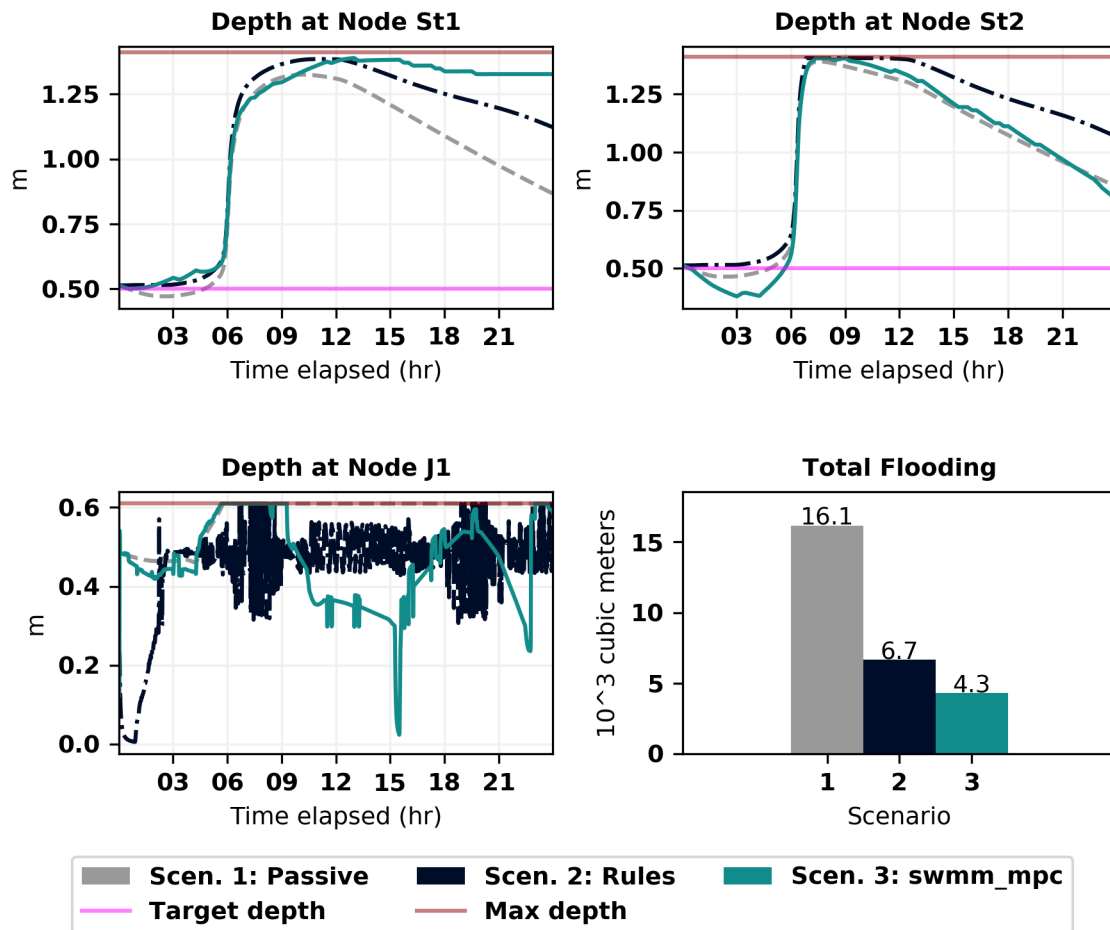


Figure 12: Depth and flooding at system nodes for Case C

to be implemented.

The fastest wall-clock time using the PC was 214.7 minutes using eight computational cores (the maximum available). Therefore the time required to find the best control policy at each control time step was 2.2 minutes. In this case, the PC’s computing power was sufficient (2.2 minutes per time step compared to 15 minute time step). For the HPC, the best case scenario was a wall-clock time of 34.8 minutes (0.36 minutes per time step) using the maximum of 28 computational cores. Although the minimum wall-clock time was achieved using all 28 cores on the HPC, the improvement in wall-clock time when increasing the number of cores past 16 was minimal. This is a relevant consideration when using a shared HPC resource where requesting more computational cores likely corresponds to a longer wait in the job queue.

The wall-clock times using GCP were much lower than the PC or HPC. In the best case, 32 vCPUs and 120 GB of RAM were used for a wall-clock time of 25.6 minutes (0.27 minutes per time step). This is a 1.3x speed-up compared to the fastest run using the HPC and almost a 8x speed-up compared to the fastest PC run. The financial cost of renting this machine was \$1.71 per hour. The GCP hardware is newer which may explain why the wall-clock time is lower even when the same number of computational cores was used.

3.3 Practical considerations

3.3.1 Computational costs

The execution times for our example use case were viable, however, a more complex model, a smaller control time step, or different EA parameters (more generations or individuals per generation) would increase the execution time. For example, the demonstration SWMM5 model required only one second or less to execute. The `swmm_mpc` workflow executed that model thousands of times. In our cases, the model was run more than 70,000 times (24-hour simulation x 4 time steps per hour x 8 generations per time step x approx. 80 individuals per generation) therefore requiring approximately 70,000 seconds of computation time (if each model takes approx. 1 second to run). If a more sophisticated SWMM5 model instance




Figure 13: Wall-clock run times for Case C with varying number of computational cores using a PC and a high performance computer.

were used, the execution time would be much higher. For example, in related research we are using a more complex model of the stormwater infrastructure for a neighborhood in Norfolk, Virginia that requires close to 60 seconds to execute for a 24-hour simulation time period. The wall-clock time for `swmm_mpc` for this more complex model would increase by around a factor of 60 compared to the simple cases demonstrated here. Assuming a linear increase, the wall-clock time would be 132 minutes/time step using the PC, thus rendering it unfeasible for running on a PC. Again assuming linear scaling, using 32 cores on GCP, the same simulation would execute at a rate of 16.7 minutes/time step, just over the 15 minute times step cutoff. To be feasible with this setup, fewer generations or fewer individuals per generation would have to be used to reduce the number of model runs per time step.

Another factor to consider for practical use of `swmm_mpc` is the control horizon and the number of control structures whose policies will be found using `swmm_mpc`. These two parameters determine the size of the overall control policy and therefore the solution space that the EA will be searching. In our example use case, the control policy was a vector of 24 bits. Therefore, there were 2^{24} possible solutions. This solution space, already large, would double if the control time step were 7.5 minutes instead of 15, or if the control horizon were

two hours instead of one. A larger solution space would result in a larger computation time to reach an effective solution.

Given the computational cost of the current `swmm_mpc` approach, the required complexity (and thus the wall-clock time) of the process model is an important consideration. A scenario with a simple process model (approx. 1 second wall-clock time) can be feasibly executed with just a PC, as shown in the system demonstration. However, a simple model may not represent complex urban drainage systems well enough to produce an effective control policy. Defining what level of detail is sufficient in the process model may be difficult, however, as it may depend on the objective of the modeling, a certain storm event, or the system itself. There is a trade-off among 1) model complexity, 2) model run time, and 3) the model’s ability to effectively simulate the relevant parts of the system. This trade-off is very relevant to the use of process models in a receding control horizon approach such as MPC and needs further research.

If a more complex model is needed, municipalities or others needing cloud-based resources to run `swmm_mpc` must consider the financial cost of renting a machine. Using the 32-core machine on GCP, the cost of finding the control policy for the 24-hour time span in Case C in our system demonstration was very low, \$0.72. This was, however, for one simple case. The cost would be higher with more complex scenarios such as a more complex model, a shorter time step, or more controls. If we assume the use of a more complex model, which takes 60 seconds to run, would increase running time by a factor of 60, that would also increase the cost by a factor of 60 to \$43.78. However, the system would only run if there were a storm in the forecast.

Given the computational cost of running the evolutionary algorithm, other, more efficient alternatives should be explored in future research. One possible alternative that is reinforcement learning (Kaelbling et al., 1996). This approach may be able to converge to a solution more quickly than an evolutionary algorithm and thus reduce run times. Another future improvement could be adding a penalty to changing actuator states and/or using another dynamic optimizer to have a less erratic behavior in the actuators.

3.3.2 Data and modeling uncertainty

The current design of `swmm_mpc` does not take into account the uncertainties in the system states, the forecast data and the process model. Because the SWMM5 engine is used to simulate both the urban drainage system and the process model, the process model assumes 1) perfect knowledge of the urban drainage system states, 2) perfect knowledge of future disturbances, and 3) perfect modeling of the urban drainage system. In a real implementation, there would be significant uncertainties in each of these aspects. In a real implementation, knowledge of the system states is available only from a limited number of sensors in the system. This data, limited in spatial and temporal resolution, would need estimated using a state estimator to set all the states in the system as is done in power systems (Abur and Expósito, 2004). More work will need to be done to investigate ways of incorporating sensor values to set the process model's initial conditions. Additionally, in the current case, the future disturbances (i.e., primarily rainfall) are known perfectly, when in reality, there is a large amount of uncertainty involved with forecasting such disturbances (see for example, Hong and Pai (2007), Valverde Ramírez et al. (2005), and Bellon and Austin (1984) regarding uncertainty in forecasting rainfall).

In addition to data uncertainties seen in reality, `swmm_mpc` does not currently consider gaps between the *simulated* behavior through the SWMM5 process model and the *actual* behavior of the urban drainage system, but assumes that simulation and reality are the same. In actuality, the gap between simulation and reality in urban drainage systems can be significant (see for example Mark et al. (2004)). On a related note, the ability for `swmm_mpc` to find a control policy that is effective for the urban drainage system is directly related to how well the process model represents the system. Given the simulation and data gaps seen in reality, the simulated results through policies found by `swmm_mpc` should be considered as the best case scenario and if the same policies were used in practice, any effects should be expected to be seen to a lesser extent. Further research is needed to determine the degree to which the results from the policies implemented in reality will differ compared to the simulation results.

swmm_mpc finds only the one best control policy which minimizes a single cost value defined by a set of weights for various objectives. This is also known as the weighted sum method for multi-objective optimization (Marler and Arora, 2010). This method, has some disadvantages including the inability to explore the solutions that have similar performance. The set of solutions that have similar overall performances with trade-offs between the different objectives (i.e., in our case reducing flooding and maintaining target water levels) is called the Pareto optimal front (Das and Dennis, 1997). Future work could be done in swmm_mpc to support the exploration of Pareto optimal solutions and the trade-offs between objectives.

4 Conclusions

A free and open-source software package, swmm_mpc, was developed which computes a control policy for controls within an urban drainage system model. The widely-used United States Environmental Protection Agency Stormwater Management Model Version 5 (SWMM5) is used to simulate the urban drainage system and the process model. A third-party Python library, pyswmm, is a critical component of the swmm_mpc workflow allowing a SWMM5 model to be run step-by-step in a Python environment. An evolutionary algorithm was used to find an effective control policy at each time step. When tested using a simple SWMM5 model, the swmm_mpc software was able to produce control policies that met objectives including minimizing flooding and minimizing deviation from target water levels at certain nodes in the system.

swmm_mpc leverages parallel computing to run the computationally expensive evolutionary algorithm more quickly. The wall-clock time for a simple SWMM5 model for a 24-hour simulation was reduced from 307 minutes to 214 minutes when the computational cores on a desktop PC were increased from two to eight. The wall-clock time was reduced even further to 34.8 minutes on a 28-core high-performance computer and to 25.6 minutes on a 32-core machine rented through the Google Cloud Platform. Parallel computing makes swmm_mpc feasible for use in real-time control with complex process models.

As the average storm intensity is projected to increase, and sea levels are expected to continue to rise, cities globally and especially on the coasts, can expect more flood conditions. Active control of urban drainage systems may be one of an array of approaches that can be used to confronting these challenges. The `swmm_mpc` software we have developed can be used, built-from, and improved upon as a tool to assist decision-makers and researchers in finding effective control policies for urban drainage systems.

Acknowledgments

We acknowledge the support of awards 1735587 (Critical and Resilient Infrastructure Systems - CRISP) and 1737432 (Smart and Connected Communities - S&CC) from the United States National Science Foundation. We thank Bryant McDonnell, one of the main developers of `pyswmm`, for his assistance in adding the hotstart file functionality to OWA-SWMM5 and `pyswmm`. We also thank the consultants of the University of Virginia Graduate Writing Lab including Kelly Cunningham, its director, for their helpful feedback in preparing the manuscript.

References

- Abur, A. and Expósito, A. G. (2004). *Power System State Estimation*, volume 24 of *Power Engineering (Willis)*. CRC Press.
- Afram, A. and Janabi-Sharifi, F. (2014). Theory and applications of HVAC control systems – A review of model predictive control (MPC). *Building and Environment*, 72:343–355.
- Bellon, A. and Austin, G. (1984). The accuracy of short-term radar rainfall forecasts. *Journal of Hydrology*, 70(1-4):35–49.
- Berggren, K., Olofsson, M., Viklander, M., Svensson, G., and Gustafsson, A.-M. (2012). Hydraulic Impacts on Urban Drainage Systems due to Changes in Rainfall Caused by Climatic Change. *Journal of Hydrologic Engineering*, 17(1):92–98.
- Bonnin, G., Martin, D., Lin, B., Parzybok, T., Yekta, M., and Riley, D. (2018). POINT PRECIPITATION FREQUENCY ESTIMATES, Norfolk, Virginia, USA. https://hdsc.nws.noaa.gov/hdsc/pfds/pfds_printpage.html?lat=36.8661&lon=-76.2890&data=depth&units=english&series=pds. Online; accessed 2018-10-04.
- Burger, G., Sitzenfrei, R., Kleidorfer, M., and Rauch, W. (2014). Parallel flow routing in SWMM 5. *Environmental Modelling & Software*, 53:27–34.

- Camacho, E. F. and Bordons, C. (2007). *Model predictive control*. Springer London Limited.
- Cembrano, G., Quevedo, J., Salamero, M., Puig, V., Figueras, J., and Martí, J. (2004). Optimal control of urban drainage systems. A case study. *Control Engineering Practice*, 12(1):1–9.
- Darsono, S. and Labadie, J. W. (2007). Neural-optimal control algorithm for real-time regulation of in-line storage in combined sewer systems. *Environmental Modelling & Software*, 22(9):1349–1361.
- Das, I. and Dennis, J. E. (1997). A closer look at drawbacks of minimizing weighted sums of objectives for pareto set generation in multicriteria optimization problems. *Structural optimization*, 14(1):63–69.
- Del Re, L., Allgöwer, F., Glielmo, L., Guardiola, C., and Kolmanovsky, I. (2010). *Automotive model predictive control: models, methods and applications*, volume 402. Springer.
- Gandomi, A. H., Yang, X.-S., Talatahari, S., and Alavi, A. H. (2013). 1 - metaheuristic algorithms in modeling and optimization. In Gandomi, A. H., Yang, X.-S., Talatahari, S., and Alavi, A. H., editors, *Metaheuristic Applications in Structures and Infrastructures*, pages 1 – 24. Elsevier, Oxford.
- Gelormino, M. S. and Ricker, N. L. (1994). Model-predictive control of a combined sewer system. *International Journal of Control*, 59(3):793–816.
- Hatt, B. E., Fletcher, T. D., and Deletic, A. (2009). Hydrologic and pollutant removal performance of stormwater biofiltration systems at the field scale. *Journal of Hydrology*, 365(3-4):310–321.
- Heusch, S. and Ostrowski, M. (2011). Model Predictive Control with SWMM. *Journal of Water Management Modeling*, 19:237–247.
- Hong, W.-C. and Pai, P.-F. (2007). Potential assessment of the support vector regression technique in rainfall forecasting. *Water Resources Management*, 21(2):495–513.
- Huber, W. C., Rossman, L. A., and Dickinson, R. E. (2005). EPA storm water management model, SWMM5. *Watershed Modeling*, CRC Press, Boca Raton, FL, pages 339–361.
- Kaelbling, L. P., Littman, M. L., and Moore, A. W. (1996). Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research*, 4:237–285.
- Kerkez, B., Gruden, C., Lewis, M., Montestruque, L., Quigley, M., Wong, B., Bedig, A., Kertesz, R., Braun, T., Cadwalader, O., Poresky, A., and Pak, C. (2016). Smarter Stormwater Systems. *Environmental Science & Technology*, 50(14):7267–7273.
- Kim, I. Y. and de Weck, O. L. (2005). Adaptive weighted-sum method for bi-objective optimization: Pareto front generation. *Structural and multidisciplinary optimization*, 29(2):149–158.
- Langson, W., Chrysoschoos, I., Raković, S., and Mayne, D. (2004). Robust model predictive control using tubes. *Automatica*, 40(1):125–133.

- Lund, N. S. V., Falk, A. K. V., Borup, M., Madsen, H., and Mikkelsen, P. S. (2018). Model predictive control of urban drainage systems: A review and perspective towards smart real-time water management. *Critical Reviews in Environmental Science and Technology*, 48(3):279–339.
- Maier, H., Kapelan, Z., Kasprzyk, J., Kollat, J., Matott, L., Cunha, M., Dandy, G., Gibbs, M., Keedwell, E., Marchi, A., Ostfeld, A., Savic, D., Solomatine, D., Vrugt, J., Zecchin, A., Minsker, B., Barbour, E., Kuczera, G., Pasha, F., Castelletti, A., Giuliani, M., and Reed, P. (2014). Evolutionary algorithms and other metaheuristics in water resources: Current status, research challenges and future directions. *Environmental Modelling & Software*, 62:271–299.
- Mark, O., Weesakul, S., Apirumanekul, C., Aroonnet, S. B., and Djordjević, S. (2004). Potential and limitations of 1D modelling of urban flooding. *Journal of Hydrology*, 299(3-4):284–299.
- Marler, R. T. and Arora, J. S. (2010). The weighted sum method for multi-objective optimization: new insights. *Structural and multidisciplinary optimization*, 41(6):853–862.
- Mayne, D., Seron, M., and Raković, S. (2005). Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224.
- Meneses, E., Gaussens, M., Jakobsen, C., Mikkelsen, P., Grum, M., and Vezzaro, L. (2018). Coordinating Rule-Based and System-Wide Model Predictive Control Strategies to Reduce Storage Expansion of Combined Urban Drainage Systems: The Case Study of Lundtofte, Denmark. *Water*, 10(1):76.
- Mitchell, M., Hershner, C., Julie, H., Schatt, D., Mason, P., and Eggington, E. (2013). Recurrent Flooding Study for Tidewater Virginia. *Virginia Institute of Marine Science, Center for Coastal Resources Management*, (January).
- Mockus, V. (2012). Storm Rainfall Depth and Distribution. In *National Engineering Handbook*, chapter 4. Natural Resources Conservation Service.
- Nelder, J. A. and Mead, R. (1965). A simplex method for function minimization. *The computer journal*, 7(4):308–313.
- Neumann, J. E., Price, J., Chinowsky, P., Wright, L., Ludwig, L., Streeter, R., Jones, R., Smith, J. B., Perkins, W., Jantarasami, L., and Martinich, J. (2015). Climate change risks to US infrastructure: impacts on roads, bridges, coastal development, and urban drainage. *Climatic Change*, 131(1):97–109.
- Puig, V., Cembrano, G., Romera, J., Quevedo, J., Aznar, B., Ramón, G., and Cabot, J. (2009). Predictive optimal control of sewer networks using CORAL tool: application to Riera Blanca catchment in Barcelona. *Water Science & Technology*, 60(4):869.
- Qin, S. and Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7):733–764.
- Sadler, J. (2018). swmm_mpc demonstration model. HydroShare.
<https://www.hydroshare.org/resource/73b38d6417ac4352b9dae38a78a47d81/>.

- Schütze, M., Campisano, A., Colas, H., Schilling, W., and Vanrolleghem, P. A. (2004). Real time control of urban wastewater systems—where do we stand today? *Journal of Hydrology*, 299(3-4):335–348.
- Sweet, W. V. and Park, J. (2014). From the extreme to the mean: Acceleration and tipping points of coastal inundation from sea level rise. *Earth’s Future*, 2(12):579–600.
- Tayfur, G., Kavvas, M. L., Govindaraju, R. S., and Storm, D. E. (1993). Applicability of St. Venant Equations for Two-Dimensional Overland Flows over Rough Infiltrating Surfaces. *Journal of Hydraulic Engineering*, 119(1):51–63.
- Valverde Ramírez, M. C., de Campos Velho, H. F., and Ferreira, N. J. (2005). Artificial neural network technique for rainfall forecasting applied to the São Paulo region. *Journal of Hydrology*, 301(1-4):146–162.
- Villarreal, E. L., Semadeni-Davies, A., and Bengtsson, L. (2004). Inner city stormwater control using a combination of best management practices. *Ecological Engineering*, 22(4-5):279–298.
- Vrabie, D., Pastravanu, O., Abu-Khalaf, M., and Lewis, F. (2009). Adaptive optimal control for continuous-time linear systems based on policy iteration. *Automatica*, 45(2):477–484.
- Zimmer, A., Schmidt, A., Ostfeld, A., and Minsker, B. (2015). Evolutionary algorithm enhancement for model predictive control and real-time decision support. *Environmental Modelling & Software*, 69:330–341.
- Zimmer, A., Schmidt, A., Ostfeld, A., and Minsker, B. (2018). Reducing Combined Sewer Overflows through Model Predictive Control and Capital Investment. *Journal of Water Resources Planning and Management*, 144(2):04017091.

Appendix A: Partial model output for Case C

2018-02-06 05:30:23 (SIMULATION TIME STEP)

ngen:	neval:	avg:	min:	max:
GENERATION	NUMBER OF	AVERAGE COST	MINIMUM COST	MAXIMUM COST
NUMBER	INDIVIDUALS			
	EVALUATED			

0	120	564.909	181.363	678.178
1	72	471.969	180.365	679.177
2	66	334.797	167.372	683.19
3	70	232.48	167.372	666.987
4	73	225.56	167.372	599.112
5	83	248.04	166.372	631.004
6	81	221.303	166.372	595.068
7	73	208.68	166.372	651.105
8	77	237.176	166.372	627.172

[0, 1, 1, 1, 0, 1, 1, 1, 1, 0, 0, 0, 1, 1, 1, 1, 1, 0, 0, 1, 1, 0, 0, 1] (BEST INDIVIDUAL)

166.37201850180506 (COST FOR BEST INDIVIDUAL)

2018-02-06 05:45:23

gen	nevals	avg	min	max
-----	--------	-----	-----	-----

0	120	548.32	174.402	689.046
1	76	429.162	174.402	590.118
2	76	305.889	174.402	598.114
3	80	228.434	174.401	682.007
4	72	225.899	174.401	689.048
5	74	229.367	174.401	621.119
6	87	220.51	174.401	586.079
7	77	227.35	174.401	599.137
8	78	212.65	174.401	688.047

[0, 1, 1, 1, 1, 1, 1, 1, 1, 0, 0, 0, 1, 1, 0, 1, 1, 0, 0, 0, 0, 0, 0, 1]

174.4006921317475

2018-02-06 06:00:23

gen	nevals	avg	min	max
-----	--------	-----	-----	-----

0	120	597.805	272.441	753.251
1	76	494.962	271.443	665.196
2	64	378.9	271.441	764.146
3	78	338.357	270.439	762.128
4	77	335.544	270.438	762.134
5	67	318.02	270.438	741.26
6	65	313.463	269.442	762.142
7	83	319.393	269.442	750.231
8	66	294.18	269.442	659.137

[1, 1, 1, 1, 0, 1, 1, 1, 1, 0, 0, 0, 1, 1, 1, 1, 1, 0, 1, 1, 1, 0, 0, 1]

269.44167254174397

2018-02-06 06:15:23

gen	nevals	avg	min	max
-----	--------	-----	-----	-----

0	120	564.68	185.448	696.096
1	78	466.118	184.447	675.068
2	69	367.825	184.447	614.188
3	79	259.746	184.447	624.176
4	60	211.748	184.447	603.112
5	73	232.349	184.447	642.16
6	84	264.941	184.444	702.072
7	64	222.8	184.444	599.161
8	54	231.009	184.444	636.161

[1, 0, 1, 1, 1, 1, 0, 1, 0, 0, 0, 0, 1, 1, 1, 1, 1, 0, 0, 0, 1, 0, 0, 1]
184.44445484477893

2018-02-06 06:30:23

gen	nevals	avg	min	max
-----	--------	-----	-----	-----

0	120	499.148	88.422	585.175
1	76	397.943	88.422	587.173
2	67	240.683	83.4249	557.114
3	77	148.278	76.4327	587.171
4	79	168.824	70.4352	579.199
5	77	144.545	70.4351	564.103
6	61	132.675	69.4349	547.064
7	72	123.795	69.4322	541.054
8	60	111.08	68.4327	540.06

[1, 1, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1]
68.43274713740458

2018-02-06 06:45:23

gen	nevals	avg	min	max
-----	--------	-----	-----	-----

0	120	501.842	115.438	626.053
1	73	362.999	115.434	595.086
2	77	232.599	115.434	598.084
3	71	178.082	115.434	622.092
4	64	165.06	115.434	593.167
5	59	172.194	115.434	630.046
6	68	174.713	115.433	607.071
7	79	169.749	115.433	618.175
8	60	152.37	115.433	572.129

[1, 1, 0, 1, 0, 0, 1, 1, 1, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0, 0, 1]
115.43347531461762

2018-02-06 07:00:23

gen	nevals	avg	min	max
-----	--------	-----	-----	-----

0	120	512.675	107.443	606.166
---	-----	---------	---------	---------

1	73	396.119	103.448	595.057
2	66	229.523	103.433	587.115
3	75	147.566	103.433	543.098
4	74	134.32	103.433	576.133
5	56	146.722	103.433	546.091
6	75	163.145	103.433	603.084
7	82	161.167	103.433	600.099
8	61	154.115	103.433	583.158

[0, 0, 0, 0, 1, 1, 0, 1, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1]

103.43264931237721

2018-02-06 07:15:23

gen	nevals	avg	min	max
0	120	499.123	100.436	613.156
1	67	374.813	98.4382	575.12
2	66	279.977	97.4315	571.23
3	79	174.216	97.4315	565.137
4	62	140.936	97.4315	567.141
5	84	165.622	97.4315	569.233
6	73	161.442	97.4315	562.058
7	71	128.182	97.4315	589.085
8	70	155.043	97.4315	572.229

[0, 1, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1]

97.43146011964107

2018-02-06 07:30:23

gen	nevals	avg	min	max
0	120	470.006	78.431	549.049
1	60	343.159	61.4292	536.06
2	70	217.512	61.4268	532.053
3	62	105.544	61.4268	548.098
4	73	123.517	60.4257	524.053
5	89	124.792	60.4284	542.193
6	65	123.809	60.4284	522.114
7	86	140.979	60.4284	526.128
8	80	117.129	60.4284	547.174

[1, 0, 1, 1, 1, 0, 0, 1, 1, 0, 0, 0, 0, 1, 1, 0, 1, 0, 0, 0, 1, 0, 0, 1]

60.42565536437247

2018-02-06 07:45:23

gen	nevals	avg	min	max
0	120	439.879	201.14	540.068
1	84	307.82	57.4158	519.129
2	76	231.506	57.4158	525.028
3	73	194.483	48.4244	506.048
4	80	124.912	46.4245	512.19

5	63	102.983	46.4245	504.226
6	70	119.73	46.4245	533.08
7	74	122.065	46.4245	525.09
8	74	111.302	46.4245	529.082

[0, 1, 1, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1]

46.42454059609455

2018-02-06 08:00:23

gen	nevals	avg	min	max
0	120	422.163	44.4102	521.154
1	67	308.373	36.4165	501.066
2	68	168.429	36.4165	493.068
3	75	97.0566	35.4167	486.044
4	81	102.572	35.4167	503.086
5	74	93.64	34.4206	490.196
6	78	88.0039	34.4206	509.173
7	74	81.4532	34.4195	524.153
8	71	81.02	34.4195	506.064

[0, 1, 1, 0, 0, 1, 0, 0, 1, 0, 0, 0, 1, 1, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]

34.41949582463466

2018-02-06 08:15:23

gen	nevals	avg	min	max
0	120	403.599	34.4007	487.029
1	76	274.299	25.4156	492.031
2	71	138.518	21.4166	483.061
3	75	68.6358	20.4164	459.132
4	66	76.4944	20.4159	480.098
5	61	75.4074	19.417	458.072
6	83	64.693	19.417	477.069
7	70	53.7713	19.417	471.097
8	83	59.6452	19.417	470.055

[0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]

19.41703181336161

2018-02-06 08:30:23

gen	nevals	avg	min	max
0	120	388.682	24.4079	484.176
1	74	258.952	19.4132	470.036
2	73	155.44	19.4132	460.045
3	82	80.9105	17.4134	462.079
4	64	68.3767	17.4134	467.127
5	79	80.5151	17.4134	493.154
6	61	70.7222	17.4134	490.161
7	78	82.9081	17.4134	475.086
8	67	79.3359	17.4134	469.048

[0, 0, 1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]
17.41342510775862

2018-02-06 08:45:23

gen	nevals	avg	min	max
-----	--------	-----	-----	-----

0	120	389.01	23.393	478.175
1	58	245.894	23.393	452.051
2	71	128.923	23.393	447.043
3	75	79.3091	23.393	450.043
4	67	64.3669	23.393	448.057
5	62	55.7484	23.393	434.062
6	71	65.7564	23.393	453.05
7	68	75.5935	23.393	454.105
8	70	69.6542	23.393	468.069

[0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 1, 0, 0, 1, 0, 0, 1]
23.392963855421687

2018-02-06 09:00:23

gen	nevals	avg	min	max
-----	--------	-----	-----	-----

0	120	375.363	37.3593	465.164
1	68	264.254	36.3588	441.062
2	77	138.681	24.3764	434.05
3	67	67.0033	15.3868	436.049
4	80	62.5066	15.3868	435.037
5	72	57.3331	15.3868	477.131
6	48	49.9949	15.3868	434.05
7	77	53.1511	15.3868	442.038
8	66	68.01	15.3868	446.082

[0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]
15.386766146993319

2018-02-06 09:15:23

gen	nevals	avg	min	max
-----	--------	-----	-----	-----

0	120	345.839	10.3741	452.131
1	62	221.061	3.40155	431.087
2	72	126.514	3.40155	442.044
3	73	75.1882	3.40155	442.153
4	72	27.9766	3.38119	441.145
5	88	28.6825	3.38119	429.059
6	72	51.2197	3.38119	422.098
7	77	28.9124	3.38119	425.036
8	67	53.9582	3.38119	428.023

[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]
3.381190826727067

2018-02-06 09:30:23

gen	nevals	avg	min	max
0	120	337.09	2.37592	431.106
1	72	181.769	2.37592	417.048
2	84	75.5149	2.37592	428.07
3	74	35.3089	2.37592	403.051
4	79	46.0723	2.37592	416.038
5	72	34.4716	2.37592	399.068
6	76	49.5477	2.37592	420.091
7	82	47.3834	2.37592	417.031
8	86	53.2014	2.37592	412.063

[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]

2.375915898617511

2018-02-06 09:45:24

gen	nevals	avg	min	max
0	120	337.087	19.3644	434.13
1	88	214.733	12.3533	408.041
2	70	96.7725	9.35952	438.134
3	70	50.5423	6.38346	420.071
4	65	43.8474	2.36987	400.052
5	78	58.5564	2.36987	440.124
6	72	44.3703	2.36987	402.043
7	79	45.0036	2.36987	404.048
8	85	46.8888	2.36987	407.023

[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]

2.3698710433763193

2018-02-06 10:00:24

gen	nevals	avg	min	max
0	120	335.603	21.3381	420.171
1	88	232.447	18.3543	418.114
2	76	136.065	8.37324	425.125
3	77	65.8346	8.37324	399.039
4	74	49.8831	2.3832	390.046
5	72	50.8198	2.36244	398.032
6	76	38.3532	2.36244	394.052
7	64	39.8222	2.36244	398.027
8	81	41.8216	2.36244	401.043

[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]

2.3624373508353225

2018-02-06 10:15:24

gen	nevals	avg	min	max
0	120	321.095	4.34934	420.171
1	77	190.977	4.34934	390.029
2	74	75.5242	2.35571	394.026

3	86	44.4021	2.35571	397.067
4	73	41.2572	2.35571	402.065
5	68	34.4718	2.35571	389.054
6	77	42.303	2.35571	380.047
7	67	52.4357	2.35571	397.054
8	84	41.7137	2.35571	413.099

[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]

2.3557102065613607

2018-02-06 10:30:27

gen	nevals	avg	min	max
0	120	317.061	8.33737	413.098
1	76	198.53	8.33737	392.079
2	62	78.3793	8.33737	417.115
3	77	52.1804	2.34777	394.094
4	73	42.1722	2.34777	385.039
5	72	48.8391	2.34777	386.065
6	73	27.2234	2.34777	376.042
7	67	48.1501	2.34777	385.079
8	75	44.7726	2.34777	386.033

[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]

2.347769801980198

2018-02-06 10:45:27

gen	nevals	avg	min	max
0	120	307.534	18.8637	401.099
1	73	177.447	10.8688	386.03
2	57	70.9445	10.8688	385.032
3	80	56.9697	10.8688	398.084
4	69	40.556	8.86587	373.043
5	79	46.8813	8.44402	379.029
6	75	56.2565	8.44402	414.122
7	76	43.9086	2.8747	379.031
8	64	47.1566	2.8747	394.106

[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 0]

2.8746986128625474