# Learning Safe and Optimal Control Strategies for Storm Water Detention Ponds ⋆

**Martijn A. Goorden** * **Kim G. Larsen** * **Jesper E. Nielsen** **
**Thomas D. Nielsen** * **Michael R. Rasmussen** ** **Jiří Srba** *

* *Department of Computer Science, Aalborg University, Aalborg,*
*Denmark (e-mail: {mgoorden,kgl,tdn,srba}@cs.aau.dk).*
** *Department of Built Environment, Aalborg University, Aalborg,*
*Denmark, (e-mail: {jen,mrr}@build.aau.dk)*

**Abstract:** Storm water detention ponds are used to manage the discharge of rainfall runoff from urban areas to nearby streams. Their purpose is to reduce the hydraulic impact and sediment loads of the receiving waters. Detention ponds are currently designed based on static controls: the output flow of a pond is capped at a fixed value. This is not optimal with respect to the current infrastructure capacity and for some detention ponds it might even violate current regulations set by the European Water Framework Directive. We apply formal methods to synthesize (i.e., derive automatically) a safe and optimal active controller. We model the storm water detention pond, including the urban catchment area and the rain forecasts, as a hybrid Markov decision process. Subsequently, we use the tool UPPAAL STRATEGO to synthesize a control strategy minimizing the cost related to pollution (optimality) while guaranteeing no emergency overflow of the detention pond (safety). Simulation results for an existing pond show that UPPAAL STRATEGO can learn optimal strategies that prevent emergency overflows, where the current static control is not always able to prevent it. At the same time, our approach can improve sedimentation during low rain periods.

*Keywords:* Stochastic hybrid systems, Switching control, Safe and optimal control, Hydroinformatics, Storm water detention ponds

## 1. INTRODUCTION

Urbanization poses two major risks related to storm water runoff management: flooding of the urban area, and environmental impact on receiving waters from hydraulic loads and pollution. The roads, roofs and other man-made surfaces of urban areas collect the rainwater and generate runoff, which needs to be transported away from the city to receiving waters to avoid urban flooding. The urban runoff carries particulates and xenobiotics from depositions on the urban surfaces (Hvitved-Jacobsen et al., 2010). Storm water detention ponds are the most commonly used storm water management tool for avoiding or reducing the impacts of storm water runoff (Tixier et al., 2011). Negative impacts of storm water runoff include unnatural disturbances, stream bed erosion, and downstream flooding (Walsh et al., 2005).

With the growth of urban areas as well as the change in climate and its related rain events, water utility companies need to constantly construct, maintain, and upgrade detention ponds to ensure efficient and safe storm water discharge. For example, it is estimated by the Danish Water and Wastewater Association that the cost of updating the Danish storm water systems will be between 0.6 and 1.3 billion euros (DANVA, 2018). Mitigating the environmental impact has increasingly changed the discharge permits over the last 15 years, so even recently constructed storm water detention ponds do not comply with present recommendations, see Jensen et al. (2020).

The requirements for the design of storm water detention ponds are, in general, based on the maximum allowed discharge flow rate into the nearby stream or river, the probability of emergency overflow, and the concentration of pollutants in the discharged water (Mobley and Culver, 2014). These regulations are derived from the European Water Framework Directive (European Commission, 2000) and discharge permits issued by the local authorities.

Currently, the most common discharge strategies involve static flows (when there is storm water in the detention pond) into the stream without taking into account the actual capacity of the receiving stream. These strategies do not necessarily comply with the regulations set by the European Water Framework Directive, as pointed out by Schütze et al. (2004).

Existing work on the design of active control strategies for urban water systems primarily focus on sewer systems and wastewater treatment plants, which do include detention ponds as a subsystem, see for example Schütze et al. (2004); Haverkort et al. (2010); Hoppe et al. (2011); Sun et al. (2017). Controlled discharge from detention
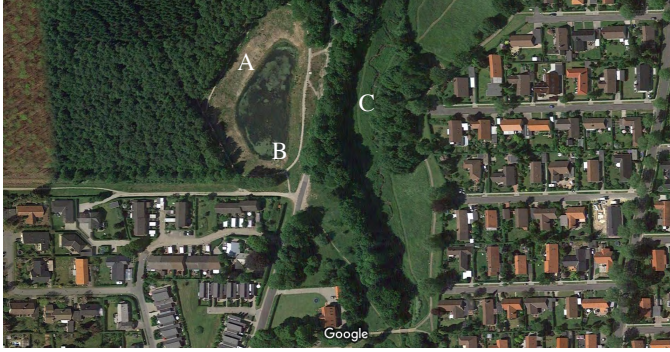
Fig. 1. Satellite image of the Vilhelmsborg Skov pond south of Aarhus, Denmark. A is the detention pond, where the outlet is located at B. Storm water is discharged in the stream labeled with C, which runs from the south to the north. Image from Google Maps.

ponds has also been studied. In Muschalla et al. (2014), a real-time controller is designed improving the efficiency of particle removal. Further improvements are presented in Gaborit et al. (2013, 2016), where off-line strategies take weather forecasts into account. However, these works focus primarily on particle removal efficiency as the objective, as most of the pollution is bound strongly to organic and inorganic particles. Furthermore, the underlying rule-based control strategies are manually derived and the discharge output can only be changed infrequently, like once a day.

In the field of cyber-physical systems, there is considerable research in deploying formal methods for verifying or even synthesizing controllers for stochastic hybrid systems. For example, the tool UPPAAL STRATEGO (David et al., 2015) is able to synthesize safe and near-optimal controllers by combining model checking and reinforcement learning. Case studies using UPPAAL STRATEGO include battery aware scheduling problems (Wognsen et al., 2015), adaptive cruise control (Larsen et al., 2015), and floor heating (Larsen et al., 2016).

In this paper, we deploy formal methods to automatically synthesize active control strategies for storm water detention ponds, where we optimize for particle sedimentation while also guaranteeing safety with respect to emergency overflows. As basis for the control synthesis, we model the storm water detention pond as a hybrid Markov decision process utilizing a combination of differential equations and (stochastic) timed automata. The synthesized controller is able to periodically change the discharge flow, thus allowing a more rapid and precise response to uncertain weather events. Simulation experiments of a real-world detention pond show that the synthesized strategies are better at utilizing the pond's capacity compared to the current static control while also honoring the stated safety requirements.

## 2. STORM WATER DETENTION PONDS

Storm water detention ponds collect storm water from urban areas, like streets, roofs, parking lots, and recreational parks. When it is raining or snow is melting, two main risks arise. First, the runoff flow in urban areas can exceed the capacity of nearby streams. Second, urban area particles
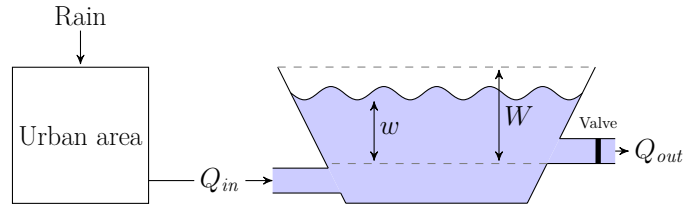


Fig. 2. Overview of the storm water detention pond

collected by the storm water can pollute the ecosystem of the stream and downstream waterbodies. Storm water detention ponds aim to reduce or avoid the impact of both risks.

An example of such a pond is shown in Figure 1. The satellite image shows the Vilhelmsborg Skov pond south of Aarhus, Denmark. Labeled by A is the storm water detention pond, partially filled with water. It collects the water from from the neighborhoods south of it through the storm water sewer system next to the roads. The pond's outlet is indicated by B, which connects to the stream labeled with C. This stream runs from south to north and discharges the water from the neighborhood.

Storm water detention ponds can be characterized by being either a wet or a dry detention pond. A wet detention pond always has a minimal amount of water in it, while a dry one can empty completely (hence the names). In our case study, we focus on a wet detention pond.

Currently, storm water detention ponds are designed with static outlet flow regulator creating a capped outlet flow into the stream. The capacity of the stream (reflected in the issued permits) dictates the maximum outlet flow of the pond. This capped outlet flow determines, together with other criteria like the urban catchment size and the emergency overflow risk, the size of the pond.

Recent research has focused on the design of energy efficient dynamic outlet flow regulator [1]. Having these flow regulators allows utility companies to incorporate active (or real-time) control into their design of the storm water detention ponds. This has the potential to enable more efficient detention pond designs and reduce the negative effect of the two aforementioned main risks.

Figure 2 shows a schematic overview of the wet storm water detention pond. Rain water falls into an urban area, like a neighborhood, university campus area, or a highway, and is transported to a nearby pond. Rain water enters the pond through inlet $Q_{in}$ and exits it through outlet $Q_{out}$ into a nearby stream or river. A valve in the outlet limits the outflow. Due to the positioning of the outlet pipe, there is a permanent water level in the pond (indicated with the lower horizontal dashed line in the figure). The variation of the water level above this permanent water level is indicated by $w$. When the maximum water level, indicated by $W$, is reached, the pond will overflow.

We neglect the rain falling directly on the pond surface, the water evaporating into the atmosphere, the water infiltration into the pond from the bottom and the sides, and the water leaking through the bottom and sides. It is

---

[1] See project webpage at `https://www.danva.dk/viden/vudp/projektuddelinger/relevand/` (in Danish).

*Rain*

*Controller*
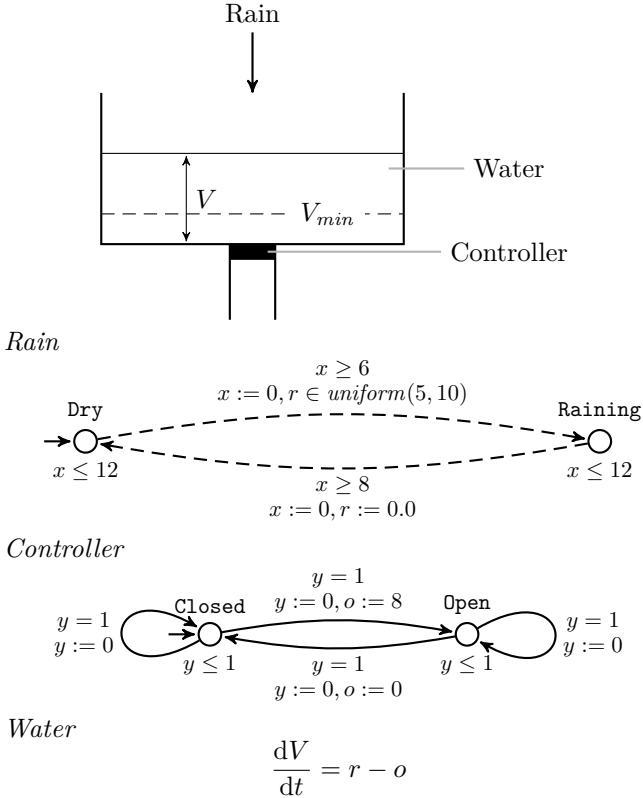
*Water*

$$\frac{\mathrm{d}V}{\mathrm{d}t} = r - o$$

Fig. 3. A small HMDP example.

shown in Thomsen et al. (2019) that these water flows are negligible compared to the storm water flow.

## 3. PRELIMINARIES

We use the mathematical modeling framework of hybrid Markov decision process (HMDP), adapted from Ashok et al. (2019); Larsen et al. (2016). The formal definitions can be found in the full conference paper (Goorden et al., 2021a). Here we will introduce HMDP with an example.

Figure 3 shows a small HMDP example. Rain can fall into a storage tank, from which water can be drained with a controlled valve. The uncontrollable environment is modeled with the (stochastic) timed automaton *Rain*. The two locations, depicted with circles, represent the two modes Dry and Raining. The clock variable $x$ keeps track of the duration of modes. The model indicates that the dry mode has a duration between 6 and 12 time units, while the raining mode has a duration between 8 and 12 time units, both uniformly distributed. Once an uncontrollable edge is taken from Dry to Raining (indicated by the dashed arrow), the rain intensity $r$ is chosen uniformly between 5 and 10 volume unit per time unit. The initial location is indicated by the small incoming arrow at location Dry. The initial value of clock variables is assumed to be 0 when not depicted.

The timed automaton *Controller* models the controllable valve, which is either in control mode Closed or Open. The solid edges indicate controllable actions. Clock variable $y$ keeps track of the control period duration, where the control period is set to 1 (see guards $y = 1$ on the edges). When switching to the Closed mode, the output flow $o$ is

set to 0 volume units per time unit, while switching to the Open mode it is set to 8 volume units per time units.

Finally, the *Water* model describes the evolution of the water volume $V$ over time with a simple differential equation: the volume change is the difference between the water inflow $r$ and the water outflow $o$.

### 3.1 Strategies for HMDP

For a given HMDP, a memoryless and possibly non-deterministic *strategy* $\sigma$ determines which of the control modes can be used in the next period. A strategy is called *deterministic* if exactly one control mode is permitted in each configuration.

A strategy $\sigma$ is called *safe* with respect to a set of states if all states encountered are within the safe set . A safe strategy is called *maximally permissive* if for each state it returns the largest set of possible actions (David et al., 2015). For the example, the safety objective could be maintaining a minimal water level $V_{min}$.

The optimality of a strategy can be evaluated for the HMDP $\mathcal{M}$ with a given optimization variable. Let $H \in \mathbb{R}_{\geq 0}$ be a given time-horizon and $D$ a random variable on finite runs, then $\mathbb{E}_{\sigma,H}^{\mathcal{M},\gamma}(D) \in \mathbb{R}_{\geq 0}$ is the expected value of $D$ on the space of runs of HMDP under strategy $\sigma$ of length $H$ starting in state $\gamma$. For the example, the optimization objective is to minimize the expected average (accumulated) water level.

### 3.2 UPPAAL STRATEGO

We use the modeling tool UPPAAL STRATEGO (David et al., 2015) for control synthesis. It integrates UPPAAL with the two branches UPPAAL SMC (Bulychev et al., 2012) (statistical model checking for stochastic hybrid systems) and UPPAAL TIGA (Behrmann et al., 2007) (synthesis for timed games). Therefore, UPPAAL is able to synthesize safe *and* optimal strategies. To synthesize a safe and optimal strategy $\sigma_{opt}$, UPPAAL STRATEGO first abstracts the HMDP into a 2-player timed game, ignoring all stochasticity. A safe strategy $\sigma_{safe}$ is afterwards synthesized for this timed game. A simplified version of timed computational tree logic (TCTL) (Behrmann et al., 2007) is used to formulate the safety specification. Subsequently, reinforcement learning is used to obtain an optimal sub-strategy $\sigma_{opt}$ based on the HMDP under supervision of the safe strategy and the given random optimization variable (David et al., 2015).

Several learning algorithms are at the modelers disposal in UPPAAL STRATEGO. Recently, in Jaeger et al. (2019) Q-learning and M-learning were introduced. With Q-learning, sample runs are drawn from the HMDP model and are used afterwards to calculate the so-called Q-values. These values are refined into a new strategy and the previous step is repeated with this new strategy until some termination criteria is met. M-learning works similar to Q-learning, except that the HMDP model is now used to approximate the transition and cost functions, which are used to calculate the Q-values instead of sample runs. To efficiently cope with continuous state spaces,
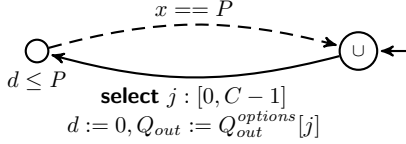
Fig. 4. The model of the controller. The **select** statement is a simplification of having a set of edges with one for each value in the select range.

UPPAAL STRATEGO uses online partition refinement techniques.

## 4. MODELING

Our model of a detention pond consists of the components Pond, Controller, Rain, and UrbanCatchment. The model can be downloaded from Goorden et al. (2021b).

*Pond* The dynamics of the water level $w$ in the pond can be derived using the mass balance equation, see Thomsen et al. (2019). Assuming constant density of water, this translates into a volume balance equation. Using Figure 2, we see that the difference in inflow and outflow contribute to the change in water inside the pond. Therefore,

$$Q_{in} - Q_{out} = \frac{dV}{dt}, \qquad (1)$$

where $Q_{in}$ is the water inlet flow from the urban drainage system, $Q_{out}$ is the water outlet flow into a nearby stream, and $V$ the water volume of the pond above the permanent water level. The outlet flow is assumed to more or less constant and equal to the discharge permission, but in reality, there will be non-linear relationship to the water level. This is however simplified in this model.

The change in volume over time can also be expressed using the geometry of the pond:

$$\frac{dV}{dt} = \frac{d(wA_p(w))}{dt}, \qquad (2)$$

where $A_p(w)$ is the pond surface area at height $w$. Equations (1) and (2) together describe the dynamics of the pond's water level under 'normal' circumstances.

There are two boundary cases that need to be taken into account. The first case is when the outflow is larger than the inflow and the water level reaches the permanent water level. The second case is when the inflow is larger than the outflow and the water level reaches the maximum height $W$ of the pond, which results in an emergency overflow. In both cases, the water level $w$ remains stationary.

Now, Equation (1) can be reformulated taking these boundary cases into account:

$$\frac{dV}{dt} = \begin{cases} 0 & \text{if } Q_{out} \geq Q_{in} \wedge w = 0, \\ 0 & \text{if } Q_{in} \geq Q_{out} \wedge w = W, \\ Q_{in} - Q_{out} & \text{otherwise.} \end{cases} \qquad (3)$$

*Controller* The controller is able to change the size of the pond outlet valve periodically. Figure 4 shows the model of the controller.

It starts in the urgent location on the right from where it choses with a controllable action one of the $C$ control options. The actual output $Q_{out}$ is set to one of the
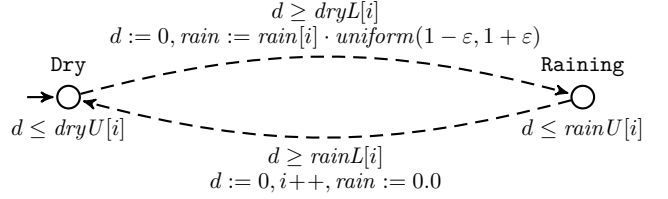


Fig. 5. The model generating rainfall.

predefined constant outputs for each mode $Q_{out}^{options}$ and the clock $d$, measuring the duration of the current control period, is reset to 0. Note that $d$ is local to the controller and does not interfere with clock $d$ from the Rain model.

In the left location, the controller waits until the period with duration $P$ is over, indicated with invariant $d \leq P$. When the controller has waited for $P$ time units, it goes to the right urgent location and above process is repeated for the next control period.

*Rain* Figure 5 shows the rain model including its uncertainty. It generates the uncontrollable input to the system. The rain profile is modeled as alternating dry and raining intervals, each modeled with a location. For each interval period, indicated with $i$, the duration of the dry (raining) period is bounded between $dryL[i]$ ($rainL[i]$) and $dryU[i]$ ($rainU[i]$), all being positive integers. Clock $d$ tracks the duration of the current interval. The actual dry or rain duration is chosen uniformly at random between $dryL[i]$ and $dryU[i]$ or $rainL[i]$ and $rainU[i]$, respectively.

When it is raining in the $i$-th interval, the actual rain intensity $rain \in \mathbb{R}$ used as input for the UrbanCatchment model is chosen uniformly random between $rain[i] \cdot (1 - \varepsilon)$ and $rain[i] \cdot (1 + \varepsilon)$, where $\varepsilon$ is a fixed uncertainty factor. Within an interval, the rain intensity is constant.

The concrete values for $dryL[i], dryU[i], rainL[i], rainU[i]$, and $rain[i]$ are derived from historical rain data from the Danish Meteorological Institute (2020).

*UrbanCatchment* The model of the urban catchment area can be found in the full version of this paper (Goorden et al., 2021a).

## 5. CONTROLLER SYNTHESIS

### 5.1 Problem definition

The main objective of the controller is to ensure a safe operation of the storm water detention pond. In this context, safety is defined as preventing the pond from overflowing. In case of an overflow event, the water discharge in the nearby stream or river is temporarily much higher than normal. This excessive discharge might have environmental impacts or cause downstream flooding.

We measure overflow with a continuous variable $o$ that represents the accumulated overflow duration. Formally,

$$\frac{do}{dt} = \begin{cases} 1 & \text{if } w = W, \\ 0 & \text{if } w < W. \end{cases} \qquad (4)$$

The secondary objective is to capture as much urban area particles as possible from the storm water. This is done

to prevent contamination of the nearby stream or river. Particles are captured through particle sedimentation onto the pond's floor surface. Hence, the more water in the pond, the more time particles have to be deposited on the floor.

In the model, we associate a cost $c$ to the ability of particle sedimentation. A linear cost function is used such that higher water levels, related to higher possibilities for particle sedimentation, result in lower cost. Formally,

$$\frac{\mathrm{d}c}{\mathrm{d}t} = 1 - \frac{w}{W}. \tag{5}$$

Therefore, $c$ represents the accumulated cost, where a cost of 1 per time unit is associated with $w$ being the permanent water level and a cost of 0 with $w$ being at the maximum height $W$.

Now, the controller synthesis problem is formulated as follows. Synthesize the safe strategy $\sigma_{safe}$ with respect to the safe state set $S$ specified by TCTL predicate $\phi \equiv \Box(o = 0)$[2], i.e., no overflow is encountered in all states of the runs according to the safe strategy. Subsequently, the optimal strategy is calculated with

$$\sigma_{opt} = \arg\min_{\sigma \subseteq \sigma_{safe}} \mathbb{E}^{\mathcal{M},\gamma}_{\sigma_{safe},H}(c), \tag{6}$$

where the particle sedimentation cost is minimized while adhering to the synthesized safe strategy $\sigma_{safe}$. Uppaal Stratego is used to synthesize controllers for this problem .

### 5.2 Experimental results

We calibrated our model to the Vilhelmsborg Skov pond south of Aarhus, Denmark. It has an urban catchment area of $A_{uc} = 0.59$ ha, a permitted discharge of 95 L/s, and an average pond area $A_p = 5,572$ m$^2$ (data from Thomsen et al. (2019)). We estimated the maximum water level to $W = 300$ cm.

Historical rain data for the period September 5 - September 7 2019 are used, obtained from the Danish Meteorological Institute (2020). For each rain event, we averaged the rain intensity. Subsequently, an uncertainty of $\varepsilon = 10\%$ is added to the observed interval durations and rain intensities to obtain a weather forecast. In this period, five rain events occurred with varying lengths and intensities. Full rain details can be found in Goorden et al. (2021a).

Figure 6 shows the results of five simulated runs in Uppaal Stratego with an initial water height $w = 100$ cm and the current static control strategy, i.e., the number of control modes $C$ for the valve is 1. In blue and solid lines the water level is plotted, in black and dotted lines the rain. We observe that one of the five runs eventually results in an emergency overflow of the pond around 2700 minutes (9 pm on September 6). This is also confirmed by analyzing the expected value of $o$: $\mathbb{E}^{\mathcal{M},w=100}_{\sigma_{static},3\text{ days}}(o) = 1.8 \pm 0.4$, i.e., the pond is expected to be overflowing for 1.8 minutes.

An actively controlled valve can have three different modes: small, medium, and large. We set the medium
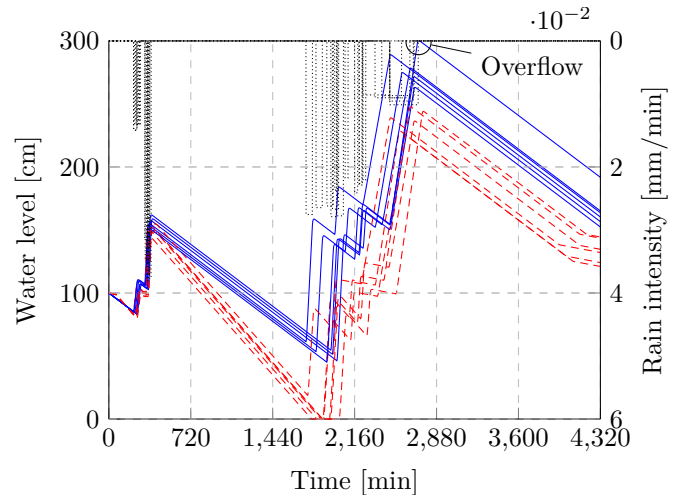
Fig. 6. Five simulations of the model with the current static control and active control. Blue and solid lines indicate the water level in the pond for static control and red and dashed lines the water level for active control. The rain fall is depicted in black, dotted, and in reverse y-axis. For one run, overflow occurs around $t = 2700$ min, as the maximum water level is $W = 300$ cm. As a quick reference, 720 minutes is a half day, and the total time scale is three days.

setting to the current static output flow capacity of 95 L/s. The low setting is 0.25 times medium and high 1.5 times medium. Due to power constraints, the valve can only change once every hour, so $P = 60$ min. We use Q-learning to synthesize strategies .

With these control modes, we synthesized an optimal controller using Equation (6). Figure 6 shows five runs in Uppaal Stratego of the model using the synthesized optimal controller (depicted in red and dashed). As can be seen from the figure, in order to ensure safety, the controller keeps the water level in the pond lower than the static controller. For some runs the pond water level even reaches the permanent water level, i.e., it cannot go lower.

Yet, having a safe strategy comes with higher cost for particle sedimentation. For static control, the expected cost is $\mathbb{E}^{\mathcal{M},w=100}_{\sigma_{static},3\text{ days}}(c) = 2026 \pm 8$, while for the optimal control it is $\mathbb{E}^{\mathcal{M},w=100}_{\sigma_{opt}3\text{ days}}(c) = 2480 \pm 8$. This is an increase of 22%, but accepted as our primary aim is to avoid flooding.

Finally, Uppaal Stratego can also report when no optimal and safe strategy can be synthesized. For example, if the initial water level is $w = 150$ cm, the query for Equation (6) cannot be satisfied. This means that no safe strategy can be found against the uncontrollable opponent (in this case the weather forecast) such that the pond will never overflow. This can be useful in predicting when emergency overflows will occur, acting as a warning system for the pond's operational organization, such that they can take additional measures at the pond in question, like maximizing the discharge flow, or in the potentially affected area downstream.

More experimental results can be found in the full conference paper (Goorden et al., 2021a).

## 6. CONCLUSION AND FUTURE CHALLENGES

We applied formal controller synthesis to automatically derive controllers for storm water detention ponds where the water discharge into the nearby stream can be regulated. We showed that the problem can be modeled as a hybrid Markov decision process, such that symbolic and reinforced learning techniques from Uppaal Stratego can be applied. Simulation results of an existing detention pond in Denmark shows that safe and near-optimal active controllers can be synthesized.

This first step opens several future research directions. First, weather forecasts change over time and are increasingly used in urban hydrology research (Thorndahl et al., 2017). Therefore, the presented model setup should be adapted to an on-line model-predictive control setting. Second, to increase the explainability of the synthesized strategies, it is to be investigated whether exporting strategies to decision trees, see Ashok et al. (2019), is possible. Third, it would be interesting to validate the approach with real-life data. Yet, actual data of water levels in ponds are scarce. Finally, only a single storm water detention pond is analyzed in isolation from the discharge stream. It is interesting to see whether collaborative strategies can be synthesized for a collection of detention ponds all discharging into the same stream.

## REFERENCES

Ashok, P., Křetínský, J., Larsen, K.G., Le Coënt, A., Taankvist, J.H., and Weininger, M. (2019). SOS: Safe, optimal and small strategies for hybrid markov decision processes. In *QEST*, LNCS, 147–164. doi:10.1007/978-3-030-30281-8_9.

Behrmann, G., Cougnard, A., David, A., Fleury, E., Larsen, K.G., and Lime, D. (2007). Uppaal-tiga: Time for playing games! In *CAV*, LNCS, 121–125.

Bulychev, P.E., David, A., Larsen, K.G., Mikučionis, M., Poulsen, D.B., Legay, A., and Wang, Z. (2012). Uppaal-smc: statistical model checking for priced timed automata. In *QALP*, 85, 1–16. EPTCS. doi:10.4204/EPTCS.85.1.

Danish Meteorological Institute (2020). DMI homepage. URL `dmi.dk`.

DANVA (2018). Administrationspraksis for regnvandsbassiner og udledningstilladelser. In Danish.

David, A., Jensen, P.G., Larsen, K.G., Mikučionis, M., and Taankvist, J.H. (2015). Uppaal stratego. In *TACAS*, LNCS, 206–211. doi:10.1007/978-3-662-46681-0_16.

European Commission (2000). Directive 2000/60/EC of the European Parliament and of the Council of 23 October 2000 Establishing a Framework for Community Action in the Field of Water Policy. Official Journal L 327.

Gaborit, E., Anctil, F., Pelletier, G., and Vanrolleghem, P.A. (2016). Exploring forecast-based management strategies for stormwater detention ponds. *Urban Water Journal*, 13(8), 841–851. doi:10.1080/1573062X.2015.1057172.

Gaborit, E., Muschalla, D., Vallet, B., Vanrolleghem, P.A., and Anctil, F. (2013). Improving the performance of stormwater detention basins by real-time control using rainfall forecasts. *Urban Water Journal*, 10(4), 230–246. doi:10.1080/1573062X.2012.726229.

Goorden, M.A., Larsen, K.G., Nielsen, J.E., Nielsen, T.D., Rasmussen, M.R., and Srba, J. (2021a). Learning safe and optimal control strategies for storm water detention ponds. URL `https://arxiv.org/abs/2104.12509`. Full ADHS'21 conference paper.

Goorden, M.A., Larsen, K.G., Nielsen, J.E., Nielsen, T.D., Rasmussen, M.R., and Srba, J. (2021b). Models for "learning safe and optimal control strategies for storm water detention ponds". Zenodo. URL `https://doi.org/10.5281/zenodo.4719450`.

Haverkort, B.R., Kuntz, M., Remke, A., Roolvink, S., and Stoelinga, M.I.A. (2010). Evaluating repair strategies for a water-treatment facility using arcade. In *DSN*, 419–424. doi:10.1109/DSN.2010.5544290.

Hoppe, H., Messmann, S., Giga, A., and Gruening, H. (2011). A real-time control strategy for separation of highly polluted storm water based on UV-Vis online measurements - from theory to operation. *Journal of Water Science and Technology*, 63(10), 2287–2293. doi:10.2166/wst.2011.164.

Hvitved-Jacobsen, T., Vollertsen, J., and Nielsen, A.H. (2010). *Urban and Highway Stormwater Pollution: Concepts and Engineering*. CRC Press. doi:10.1201/9781439826867.

Jaeger, M., Jensen, P.G., Larsen, K.G., Legay, A., Sedwards, S., and Taankvist, J.H. (2019). Teaching stratego to play ball: Optimal synthesis for continuous space MDPs. In *ATVA*, LNCS, 81–97. doi:10.1007/978-3-030-31784-3_5.

Jensen, D.M.R., Thomsen, A.T.H., Larsen, T., Egemose, S., and Mikkelsen, P.S. (2020). From EU directives to local stormwater discharge permits: A study of regulatory uncertainty and practice gaps in Denmark. *Sustainability*, 12(16), 6317. doi:10.3390/su12166317.

Larsen, K.G., Mikučionis, M., Muñiz, M., Srba, J., and Taankvist, J.H. (2016). Online and compositional learning of controllers with application to floor heating. In *TACAS*, LNCS, 244–259. doi:10.1007/978-3-662-49674-9_14.

Larsen, K.G., Mikučionis, M., and Taankvist, J.H. (2015). Safe and optimal adaptive cruise control. In *Olderog-Festschrift*, LNCS, 260–277. doi:10.1007/978-3-319-23506-6_17.

Mobley, J.T. and Culver, T.B. (2014). Design of outlet control structures for ecological detention ponds. *Journal of Water Resources Planning and Management*, 140(2), 250–257. doi:10.1061/(ASCE)WR.1943-5452.0000266.

Muschalla, D., Vallet, B., Anctil, F., Lessard, P., Pelletier, G., and Vanrolleghem, P.A. (2014). Ecohydraulic-driven real-time control of stormwater basins. *Journal of Hydrology*, 511, 82–91. doi:10.1016/j.jhydrol.2014.01.002.

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. doi:10.1016/j.jhydrol.2004.08.010.

Sun, C., Joseph-Duran, B., Maruejouls, T., Cembrano, G., Meseguer, J., Puig, V., and Litrico, X. (2017). Real-time control-oriented quality modelling in combined urban drainage networks. *IFAC-PapersOnLine*, 50(1), 3941–3946. doi:10.1016/j.ifacol.2017.08.142.

Thomsen, A.T.H., Nielsen, J.E., and Rasmussen, M.R. (2019). A simplified method for measuring the discharge from stormwater detention ponds. Submitted.

Thorndahl, S., Einfalt, T., Willems, P., Nielsen, J.E., ten Veldhuis, M.C., Arnbjerg-Nielsen, K., Rasmussen, M.R., and Molnar, P. (2017). Weather radar rainfall data in urban hydrology. *Hydrology and Earth System Sciences*, 21(3), 1359–1380. doi:10.5194/hess-21-1359-2017.

Tixier, G., Lafont, M., Grapentine, L., Rochfort, Q., and Marsalek, J. (2011). Ecological risk assessment of urban stormwater ponds: literature review and proposal of a new conceptual approach providing ecological quality goals and the associated bioassassment tools. *Journal of Ecological Indicators*, 11(6), 1497–1506. doi:10.1016/j.ecolind.2011.03.027.

Walsh, C.J., Roy, A.H., Feminella, J.W., Cottingham, P.D., Groffman, P.M., and Morgan, R.P. (2005). The urban stream syndrome: current knowledge and the search for a cure. *Journal of the North American Benthological Society*, 24(3), 706–723. doi:10.1899/04-028.1.

Wognsen, E.R., Haverkort, B.R., Jongerden, M., Hansen, R.R., and Larsen, K.G. (2015). A score function for optimizing the cycle-

life of battery-powered embedded systems. In *FORMATS*, LNCS, 305–320. doi:10.1007/978-3-319-22975-1_20.