`pipedream`: an interactive digital twin model for urban drainage networks

Matthew Bartos[a,*], Branko Kerkez[a]

[a]*Department of Civil and Environmental Engineering, University of Michigan, Ann Arbor, Michigan, USA.*

Abstract

Faced with growing water infrastructure challenges, many cities are now seeking to build *digital twins* of urban stormwater systems that combine sensor data with online models in order to better understand and control system dynamics. Towards this goal, this study presents `pipedream`—an end-to-end software toolkit for real-time modeling and state estimation in urban stormwater networks. The toolkit combines (i) a new hydrodynamic solver based on the full one-dimensional Saint-Venant equations and (ii) an implicit Kalman filtering methodology that efficiently updates system states based on observed data. Drawing on sensor data from a real-world stormwater network, we find that the state estimation toolkit is effective at both interpolating system states and forecasting future states based on current measurements. By providing a complete, real-time view of stormwater system dynamics, this toolkit will enable better evaluation of system performance, improved detection of hazards, and more robust implementation of real-time control.

*Corresponding author

Email addresses: `mdbartos@umich.edu` (Matthew Bartos), `bkerkez@umich.edu` (Branko Kerkez)

1 Introduction

In the wake of growing urban populations, aging infrastructure, and more frequent extreme weather events, many cities are struggling to manage stormwater-related challenges such as flash flooding and combined sewer overflows [1]. Engineers have traditionally responded to these challenges by expanding stormwater infrastructure and implementing best management practices [2]. However, despite the high costs of these interventions, it is often difficult to evaluate their benefits and impacts due to (i) a lack of real-time data and (ii) a limited understanding of system dynamics [1, 3]. Interventions may fail to achieve performance targets due to design oversights, insufficient maintenance, or changing hydrologic conditions [4, 5]. In certain pathological cases, measures aimed at improving flood control and urban water quality may actually worsen the problems they are intended to solve [6, 7]. In the absence of continuous monitoring, these issues may remain undetected until negative impacts have already accrued.

To ensure that stormwater systems achieve desired outcomes, water managers are now seeking to build *digital twins* of real-world networks that use embedded sensors and online models to monitor system dynamics in real-time. While these digital twins vary widely in scope and sophistication, they share the general goal of integrating sensor data, hydrodynamic models, and data assimilation techniques in order to assess system performance and determine effective control strategies. Using these technologies, operators can detect abnormal conditions within sewer networks and then dispatch maintenance crews to make repairs before damage occurs [5, 8]. Moreover, when combined with real-time control (RTC), continuous monitoring has shown impressive results in reducing combined sewer overflows [8–12], reducing operational costs [5, 8], and improving urban water quality [13–16].

While digital twins promise to address many of stormwater management's biggest pain points, adoption has been hampered by a lack of sufficient tools and theory for online hydraulic modeling and real-time data assimilation. In practice, hydraulic models are primarily executed in batch mode for the purposes of sizing pipes and evaluating infrastructure expansions [17]. Popular stormwater models are oriented towards this use case, and relatively few software packages provide support for online modeling or data assimilation. Consequently, for real-world systems that seek to implement continuous

monitoring, the underlying process model is often substantially simplified—for instance, assuming that the stormwater network behaves as a series of cascading linear reservoirs [11]. These simplifications may introduce uncertainty when interpolating or forecasting system states, which in turn may impair the performance of real-time control strategies. New tools are needed to enable online state estimation and control capabilities while at the same time ensuring that system dynamics are accurately represented.

To pave the way for continuous monitoring and control of urban drainage systems, we introduce `pipedream`—a software toolkit for building digital twins of stormwater networks. This toolkit consists of (i) an efficient hydraulic solver based on the full one-dimensional Saint-Venant equations and (ii) an implicit Kalman filtering approach that updates system states based on streaming sensor data (see Figure 1). Using the full one-dimensional Saint-Venant equations ensures that the model adequately captures the physics of a wide variety of real-world stormwater networks. Similarly, Kalman filtering enables fast assimilation of sensor data while at the same time respecting the dynamics of the physical system. `pipedream` can be run in online mode, advancing forward in sync with the real-world system and assimilating sensor measurements in real-time. Moreover, the toolkit provides a robust interface for executing dynamic controls through the use of adjustable orifices, weirs and pumps. Taken together, the software described in this paper provides an end-to-end framework for real-time monitoring and control that will enable system operators to better manage stormwater infrastructure.

2 Prior Work

Previous efforts towards interactive stormwater modeling have largely focused on developing enhancements to existing models, such as the EPA Stormwater Management Model (SWMM) [18]. One of the earliest of these efforts was MATSWMM: a Mat-Lab/Python library aimed at simulating real-time control strategies for stormwater systems [19]. The MATSWMM library provides bindings for the SWMM hydrodynamic solver, and also offers tools for implementing control strategies; however, at the time of writing, the project is no longer actively maintained. pySWMM is an actively-maintained Python library that also provides interactive bindings for SWMM, allowing users to modify system states and simulate real-time control strategies [20]. While pySWMM has been
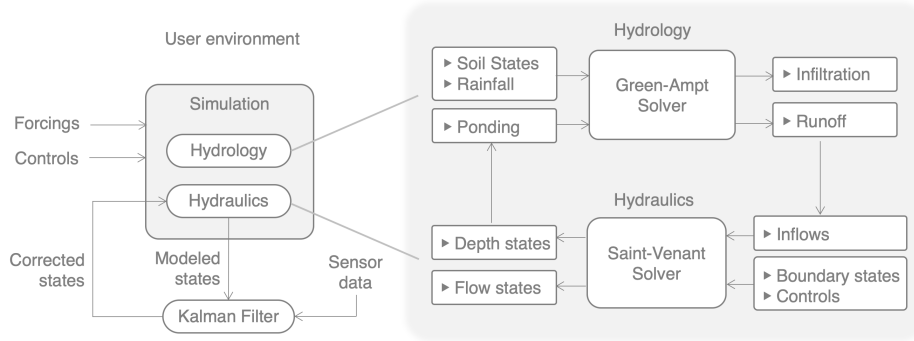
Figure 1: Software diagram of the `pipedream` toolkit. The simulation engine consists of a coupled hydraulic/hydrologic solver based on the Saint-Venant and Green-Ampt equations. The solver simulates stormwater dynamics in real-time by combining rainfall forcings, user-specified control inputs, and online sensor data which is fused into the model using a Kalman Filter.

used to great effect in simulating real-time control strategies within industry, its application to continuous monitoring has been limited by the fact that SWMM does not offer a state-space representation of system dynamics, making it difficult to apply model-based data assimilation techniques like Kalman Filtering.[1]

On the data assimilation side, a significant body of work has explored the problem of state/parameter estimation in hydrologic systems. Much of this research focuses on the problem of streamflow forecasting in large river basins [23]. These studies seek to improve streamflow forecasts by using data assimilation techniques to correct initial soil moisture conditions [24–26], snow water storage [27], and rainfall forcings [24]. A variety of data assimilation approaches have been investigated, including variational data assimilation [24], best linear unbiased estimation [25], particle filtering [27, 28], and ensemble Kalman filtering [26, 27, 29]. While most of the literature focuses on offline data assimilation, a few studies have examined data assimilation in a real-time context. Schwanenberg (2011) develop a data assimilation approach based on Ensemble Kalman filtering to assist with real-time control of large delta river systems [30]. However, their approach uses the

---

[1]Although it is possible to assimilate data using a surrogate model when the underlying process model is unavailable [21], this approach often lacks the performance guarantees of model-based data assimilation approaches [22].

kinematic wave approximation of the Saint-Venant equations, rendering it unsuitable for stormwater networks with bidirectional flow.

In the context of stormwater networks, the application of data assimilation has been more limited. In practice, efforts towards incorporating sensor data into stormwater models have mainly focused on model parameter calibration [31–35]. Online state estimation has generally been achieved through model-specific deterministic methods that are difficult to reconcile with modern state estimation theory [36, 37]. Studies that explore more formal state estimation techniques generally rely on simplified conceptual models. Hutton et al. (2014), for instance, present a method for applying a Kalman Filter to a simplified stormwater model based on a series of cascading linear reservoirs [38], while Breinholt et al. (2011, 2012) investigate the application of Extended Kalman Filtering to a lumped conceptual urban rainfall-runoff model [39, 40]. While this line of research shows the potential for data assimilation to improve our understanding of stormwater system performance, more work is needed to integrate robust data assimilation techniques with our best knowledge of system dynamics, as represented by physically-based models. To our knowledge, there is currently no fully physically-based interactive stormwater model that provides real-time data assimilation capabilities. To fill this need, we present `pipedream`—a new digital twin model for stormwater networks.

3   Methods

In this section, we describe the development of the `pipedream` toolkit, including the theory behind the hydrodynamic model, its numerical implementation, and the development of a state estimation procedure for fusing real-time sensor data. First, we develop and implement a robust hydraulic solver that facilitates data assimilation by providing a state space model of the stormwater system (Section 3.1). Next, we combine the hydraulic solver with an implicit Kalman filter to facilitate real-time assimilation of sensor data into the dynamical model (Section 3.3). Finally, drawing on sensor data from a real-world stormwater network, we test the data assimilation framework by evaluating its efficacy at interpolating and forecasting system states (Section 3.6).

5

3.1 Overview of the hydraulic solver

To enable real-time state estimation in urban drainage systems, we first develop and implement a new hydraulic solver that (i) enables real-time interactive usage, and (ii) provides a state space representation of system dynamics. The hydraulic solver developed in this study is based on an implicit staggered-grid scheme for solving the one-dimensional Saint-Venant equations in sewer/channel networks [41]. This formulation has demonstrated impressive accuracy and stability in previous applications, and variants of the scheme are used in a number of proprietary software packages, including GSSHA and SEWERCAT [42, 43]. More crucially for this study, this numerical scheme facilitates the use of robust state estimation techniques like Kalman filtering by enabling the system dynamics to be cast in the form of an implicit state space model. In this section, we discuss the basic solver formulation, and show how the model can be reinterpreted as an implicit linear time-varying (LTV) state space system. We also present several improvements to the original solver formulation that enhance model stability, enable representation of backwater effects, and extend the applicability of the model to systems with dynamic control structures.
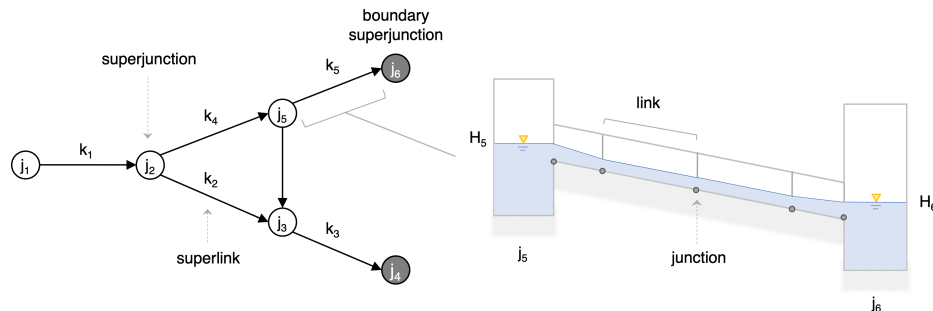

Figure 2: Example network from Ji (1998) demonstrating the fundamental computational elements used by the model [41], including superlinks/superjunctions (left), and links/junctions (right).

The hydraulic solver is a finite difference model with four distinct types of computational elements: links, junctions, superlinks, and superjunctions [41]. *Links* are finite volumes that represent sections of conduit or open channel. *Junctions* connect links together, and may represent manholes, grade changes, or simply extra computational

elements. *Superlinks* are collections of links joined end-to-end by junctions with no branching. *Superjunctions* connect one or more superlinks together, and may represent storage basins, branching locations, or invert discontinuities. To facilitate the computation of pressurized flow, each link is equipped with a Priessman slot—a fictitious narrow groove located at the crown of the link that allows the one-dimensional unsteady open-channel flow equations to be applied to surcharged pipes [44]. Figure 2 shows an example network adapted from Ji (1998) [41], with all basic elements of the model labeled.

Flows within the channel network are modeled using the one-dimensional Saint-Venant equations. This pair of nonlinear partial differential equations consists of two parts: the continuity equation (1), which describes the mass balance for a finite volume, and the momentum equation (2), which describes the force balance:

$$\frac{\partial A}{\partial t} + \frac{\partial Q}{\partial x} = q_{in} \tag{1}$$

$$\frac{\partial Q}{\partial t} + \frac{\partial}{\partial x}(Qu) + gA\left(\frac{\partial h}{\partial x} - S_0 + S_f + S_L\right) = 0 \tag{2}$$

Where Q is discharge; A is the cross-sectional area of flow; u is the average velocity; h is depth; x is distance; t is time; q_{in} is the lateral inflow per unit width; and S_o, S_f and S_L represent the channel bottom slope, friction slope and local head loss slope, respectively. Using a staggered grid formulation, the continuity equation is applied to each junction (indexed by Ik), while the momentum equation applied to each link (indexed by ik). The equations are discretized using a Backward Euler-type implicit scheme (see Sections S2 and S3 in the Supplementary Information for derivations):

$$Q_{ik}^{t+\Delta t} - Q_{i-1k}^{t+\Delta t} + \left(\frac{B_{ik}\Delta x_{ik}}{2} + \frac{B_{i-1k}\Delta x_{i-1k}}{2} + A_{s,Ik}\right) \cdot \frac{h_{Ik}^{t+\Delta t} - h_{Ik}^t}{\Delta t} = Q_{in,Ik} \tag{3}$$

$$(Q_{ik}^{t+\Delta t} - Q_{ik}^t)\frac{\Delta x_{ik}}{\Delta t} + u_{I+1k}Q_{I+1k}^{t+\Delta t} - u_{Ik}Q_{Ik}^{t+\Delta t}$$
$$+gA_{ik}(h_{I+1k}^{t+\Delta t} - h_{Ik}^{t+\Delta t}) - gA_{ik}S_{o,ik}\Delta x_{ik} + gA_{ik}(S_{f,ik} + S_{L,ik})\Delta x = 0 \tag{4}$$

Where B is the top width of flow, A_s is the junction surface area, and Q_{in} is the exogenous flow input. The boundary conditions for each superlink are supplied by the

upstream and downstream superjunction heads, assuming weir-like flow at the superlink inlet and outlet:

$$Q = CA\sqrt{2g\Delta H} \tag{5}$$

Where C is the inlet/outlet discharge coefficient, and ΔH is the difference in head between the superjunction and the adjacent superlink boundary junction.

The hydraulic model solves for all unknowns simultaneously at each time step by embedding the solutions to the Saint-Venant equations into a system of implicit linear equations wherein all unknowns are defined in terms of the unknown superjunction heads.

1. First, the discretized Saint-Venant equations are reformulated into recurrence relations that relate junction heads and link flows within each superlink.

2. The assumption of orifice-like flow between superjunctions and superlinks is used to establish boundary conditions for the superlink inlets and outlets.

3. Combining the recurrence relations together with the superlink boundary conditions, the system is reformulated as a sparse matrix equation with all unknowns expressed in terms of the unknown superjunction heads.

4. After solving for the unknown superjunction heads, the internal depths and flows within each superlink are recovered by substituting the superjunction boundary heads into the previously-developed recurrence relations.

This solution procedure affords a balance between stability and computational efficiency: while the implicit discretization scheme helps to ensure stability, the use of recurrence relations reduces the size of the solution matrix and helps to increase computation speed compared to the more conventional four-point implicit scheme [41]. However, in addition to its computational advantages, this scheme also facilitates data assimilation by enabling the solver to be cast in the form of a linear time-varying state space system.

3.2 Constructing the state space system

Through derivation of the solution matrix equation, we show that the hydraulic model is equivalent to a linear time-varying (LTV) state space system, which in turn allows for the application of powerful algorithms from signal processing and control theory, such

as feedback control, modal analysis, and—as we show in this paper—Kalman filtering. In the following section, we derive the solution matrix equation and show how it is equivalent to a LTV state space system. Detailed derivations, including the discretization of governing equations and the development of the recurrence relations, are omitted from this section and can instead be located in the Supplementary Information (SI) document.

The solution matrix equation is derived by applying the continuity equation to all superjunctions, and then using recurrence relations to express all unknown variables in terms of the superjunction heads. Focusing on a single superjunction j, the continuity equation requires that the change in storage volume over a time interval Δt equal the sum of inflows minus outflows:

$$\sum_{l=1}^{NKDj} Q_{dk_l}^{t+\Delta t} - \sum_{m=1}^{NKUj} Q_{uk_m}^{t+\Delta t} + Q_{in,j}^{t+\Delta t} = \frac{A_{sj}(H_j^{t+\Delta t} - H_j^t)}{\Delta t} \tag{6}$$

Where Q_{dk} is the discharge at the downstream end of superlink k, Q_{uk} is the discharge at the upstream end of superlink k, $Q_{in,j}$ is the external inflow into superjunction j, A_{sj} is the surface area of superjunction j, and H_j is the total head at superjunction j. NKD_j represents the number of superlinks with their downstream ends attached to superjunction j, while NKU_j represents the number of superlinks with their upstream ends attached to superjunction j (thus, the first two terms represent the sum of inflows minus outflows from all superlinks attached to superjunction j).

Using the recurrence relations developed in SI Section S4, the discharge at the upstream and downstream ends of each superlink can be described using the following linear functions of the upstream and downstream superjunction heads at time $t + \Delta t$:

$$Q_{uk}^{t+\Delta t} = \alpha_{uk} H_{juk}^{t+\Delta t} + \beta_{uk} H_{jdk}^{t+\Delta t} + \chi_{uk} \tag{7}$$

$$Q_{dk}^{t+\Delta t} = \alpha_{dk} H_{juk}^{t+\Delta t} + \beta_{dk} H_{jdk}^{t+\Delta t} + \chi_{dk} \tag{8}$$

Where H_{juk} represents the head at the superjunction upstream of superlink k, and H_{jdk} represents the head at the superjunction downstream of superlink k. The α, β and χ coefficients are functions of the depths and flows inside each superlink that incorporate the solutions to the continuity and momentum equations in both the forward and backward directions. Substituting these linearized expressions into the continuity balance for

9

superjunction j yields the linear equation:

$$\left(\frac{A_{sj}}{\Delta t} + F_{j,j}^k\right) \cdot H_j^{t+\Delta t} + \sum_{\ell=1}^{NKDj} \Phi_{j,juk_\ell}^k H_{juk_\ell}^{t+\Delta t} + \sum_{m=1}^{NKUj} \Psi_{j,jdk_m}^k H_{jdk_m}^{t+\Delta t}$$
$$= \frac{A_{sj}}{\Delta t} H_j^t + Q_{in,j}^{t+\Delta t} + G_j^k \tag{9}$$

Where:

$$F_{j,j}^k = \sum_{m=1}^{NKUj} \alpha_{uk_m} - \sum_{\ell=1}^{NKDj} \beta_{dk_\ell} \tag{10}$$

$$\Phi_{j,juk_\ell}^k = -\alpha_{dk_\ell} \tag{11}$$

$$\Psi_{j,jdk_m}^k = \beta_{uk_m} \tag{12}$$

$$G_j^k = \sum_{\ell=1}^{NKDj} \chi_{dk_\ell} - \sum_{m=1}^{NKUj} \chi_{uk_m} \tag{13}$$

Applying these equations to each superjunction j yields the following implicit state space equation:

$$\left(A_K(t) + A_S(t)\right) \cdot \mathbf{x}(t + \Delta t) = A_S(t)\mathbf{x}(t) + B(t)\mathbf{u}(t) + D_K(t) \tag{14}$$

Where $\mathbf{x}(t)$ is the state vector of superjunction heads at the current time step and $\mathbf{x}(t + \Delta t)$ is the state vector of superjunction heads at the next time step. The elements of $A_K(t)$ are defined such that $A_{K_{j,j}}(t) = F_{j,j}^k$; $A_{K_{j,juk_\ell}} = \Phi_{j,juk_\ell}^k$; and $A_{K_{j,jdk_m}} = \Psi_{j,jdk_m}^k$ (with all other elements of $A_K(t)$ equal to zero). Similarly, the elements of $D_K(t)$ are defined such that $D_{K_j} = G_j^k$. $A_S(t)$ is a diagonal matrix representing the current superjunction storage:

$$A_S(t) = \text{diag}\left(\frac{A_{sj}}{\Delta t}\right), \quad \forall j = 1, 2, ..., M \tag{15}$$

10

The exogenous input $B(t)\mathbf{u(t)}$ represents the external inflow (e.g. runoff) to each superjunction j. Assuming that each superjunction has its own independent input signal, the input transition matrix is the identity matrix:

$$B(t) = I \tag{16}$$

And the input signal is equal to the external inflow at each superjunction:

$$\mathbf{u}(t) = [Q_{in,1}, \ Q_{in,2}, \cdots Q_{in,M}]^T \tag{17}$$

To illustrate the construction of the solution matrix in concrete terms, we include here the solution matrix for the example network in Figure 2. For brevity, define $\tilde{F}_{j,j} = \frac{A_{sj}}{\Delta t} + F_{j,j}^k$ and $\tilde{G}_j = \frac{A_{sj}}{\Delta t} H_j^t + Q_{in,j} + G_j^k$. Thus, for the example network in Ji (1998) [41], the sparse matrix equation at time $t + \Delta t$ is expressed as:

$$
\begin{bmatrix}
\tilde{F}_{1,1} & \Psi_{1,2}^k & 0 & 0 & 0 & 0 \\
\Phi_{2,1}^k & \tilde{F}_{2,2} & \Psi_{2,3}^k & 0 & \Psi_{2,5}^k & 0 \\
0 & \Phi_{3,2}^k & \tilde{F}_{3,3} & \Psi_{3,4}^k & \Phi_{3,5}^k & 0 \\
0 & 0 & 0 & 1 & 0 & 0 \\
0 & \Phi_{5,2}^k & \Psi_{5,3}^k & 0 & \tilde{F}_{5,5} & \Psi_{5,6}^k \\
0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}
\begin{bmatrix}
H_1^{t+\Delta t} \\
H_2^{t+\Delta t} \\
H_3^{t+\Delta t} \\
H_4^{t+\Delta t} \\
H_5^{t+\Delta t} \\
H_6^{t+\Delta t}
\end{bmatrix}
=
\begin{bmatrix}
\tilde{G}_1 \\
\tilde{G}_2 \\
\tilde{G}_3 \\
H_{bc,4}^{t+\Delta t} \\
\tilde{G}_5 \\
H_{bc,6}^{t+\Delta t}
\end{bmatrix}
\tag{18}
$$

This sparse system is solved for the unknown left-hand vector of superjunction heads at time step $t + \Delta t$. Once the superjunction heads are known, the depths and discharges at the upstream and downstream superlink boundaries are computed. Finally, the recurrence relations are used to solve for the internal depths and flows within each superlink.

3.3   Implicit Kalman Filter

Having defined the system dynamics in terms of an implicit LTV system, an implicit Kalman Filter is developed to fuse sensor data with the dynamical model. Kalman Filtering is a recursive Bayesian estimation algorithm that (i) uses a dynamical system model to generate a prior estimate of system states, then (ii) updates this prior with observed data to produce a posterior estimate [45]. It can be shown that the Kalman Filter is the optimal linear estimator for system states when the dynamical system model

11

is perfectly accurate, and the noise is white and Gaussian with known covariance [45]. In the following section, we derive the Kalman recursion for an implicit LTV system, following the treatment provided in Skliar and Ramirez (1997) [46]. As a starting point, we may first express the internal states and observed outputs of the system in terms of a "hidden" state equation along with an observed output equation.

$$A_1(t)\mathbf{x}(t + \Delta t) = A_2(t)\mathbf{x}(t) + B(t)\mathbf{u}(t) + C(t)\mathbf{w}(t) + D(t) \tag{19}$$

$$\mathbf{z}(t + \Delta t) = H(t + \Delta t)\mathbf{x}(t + \Delta t) + \mathbf{v}(t + \Delta t) \tag{20}$$

Where $\mathbf{x}(t)$ is the $n-$dimensional state vector at the current time step, $\mathbf{u}(t)$ is an $\ell-$dimensional input vector, $\mathbf{w}(t)$ is a $p-$dimensional stochastic disturbance, $\mathbf{z}(t + \Delta t)$ is the $m-$dimensional observed state at the next time step, and $\mathbf{v}(t + \Delta t)$ is the $m-$dimensional vector of measurement noise. $A_1(t)$ and $A_2(t)$ are $(n \times n)$ state transition matrices, $B(t)$ is an $(n \times \ell)$ input transition matrix, $C(t)$ is an $(n \times p)$ stochastic disturbance transition matrix, $D(t)$ is an $(n \times 1)$ vector of constants, and $H(t + \Delta t)$ is the $(m \times n)$ measurement transition matrix.

The stochastic disturbance $\mathbf{w}(t)$, and measurement noise $\mathbf{v}(t + \Delta t)$ are assumed to be random vectors of zero-mean Gaussian white noise:

$$E[\mathbf{w}(t_2)\mathbf{w}(t_1)^T] = Q(t)\delta(t_2 - t_1) \tag{21}$$

$$E[\mathbf{v}(t_2 + \Delta t)\mathbf{v}(t_1 + \Delta t)^T] = R(t + \Delta t)\delta(t_2 - t_1) \tag{22}$$

Where $Q(t)$ is the $(p \times p)$ covariance matrix for the process noise and $R(t + \Delta t)$ is the $(m \times m)$ covariance matrix for the measurement noise. To apply a Kalman filter to a system of this form, we may first define the auxiliary state vector [46]:

$$\mathbf{y}(t + \Delta t) = A_1(t)\mathbf{x}(t + \Delta t) \tag{23}$$

An estimate of the auxiliary state vector at the next time step is produced by propagating the dynamical model forward in time:

$$\hat{\mathbf{y}}_{t+\Delta t|t} = A_2(t)\hat{\mathbf{x}}_{t|t} + B(t)\mathbf{u}(t) + D(t) \tag{24}$$

12

Next, the estimation error covariance of $\mathbf{y}(t + \Delta t)$ is computed by propagating the previous error covariance forward in time through the dynamical model:

$$P^y_{t+\Delta t|t} = A_2(t)P^x_{t|t}A_2^T(t) + C(t)Q(t)C^T(t) \tag{25}$$

The auxiliary measurement matrix H_1 is defined as:

$$H_1(t + \Delta t) = H(t + \Delta t)A_1^{-1}(t) \tag{26}$$

The optimal Kalman gain can then be computed:

$$L^y(t + \Delta t) = P_{t+\Delta t|t}H_1(t + \Delta t)\left[H_1(t + \Delta t)P^y_{t+\Delta t|t}H_1^T(t + \Delta t) + R(t + \Delta t)\right]^{-1} \tag{27}$$

The updated estimate of the auxiliary state vector is obtained by applying the Kalman gain to the forward propagation of the dynamical system model:

$$\hat{\mathbf{y}}_{t+\Delta t|t+\Delta t} = A_2(t)\hat{\mathbf{x}}_{t|t} + B(t)\mathbf{u}(t) + D(t) + L^y(t+\Delta t)\left[\mathbf{z}(t+\Delta t) - H_1(t+\Delta t)\hat{\mathbf{y}}_{t+\Delta t|t}\right] \tag{28}$$

The updated estimate of the original state vector is recovered through inversion of the left state transition matrix:

$$\hat{\mathbf{x}}_{t+\Delta t|t+\Delta t} = A_1^{-1}(t)\hat{\mathbf{y}}_{t+\Delta t|t+\Delta t} \tag{29}$$

The covariance measurement update for the auxiliary state vector is then computed as:

$$P^y_{t+\Delta t|t+\Delta t} = (I - L^y(t + \Delta t)H_1(t + \Delta t))P^y_{t+\Delta t|t} \tag{30}$$

Finally, the covariance measurement update for the original state vector is recovered through inversion of the left state transition matrix:

$$P^x_{t+\Delta t|t+\Delta t} = A_1^{-1}(t)P^y_{t+\Delta t|t+\Delta t}(A_1^{-1}(t))^T \tag{31}$$

13

After computing the updated covariance measurement, the process is repeated for the next time step, starting with equation 24 and ending with equation 31.[2]

The data fusion procedure proceeds as follows: at each time step, the hydraulic solver is advanced to determine the hydraulic heads and update the state transition matrices. Next, the Kalman recursion is applied to correct the hydraulic heads at time step $t + \Delta t$ based on observed data. Finally, the superlink boundary conditions and internal states are computed based on the updated superjunction heads. This process is repeated indefinitely (potentially in real-time) until the simulation is terminated.

3.4 Enhancements to the hydraulic solver

To enable representation of a broader array of real-world stormwater networks, we make substantial additions to the original numerical scheme. These changes allow for modeling of real-world stormwater networks that would be either impractical or impossible under the original formulation.

Control structures: To facilitate simulation and execution of real-time controls, we modify the numerical scheme to enable representation of orifices, weirs, and pumps (see Sections S9, S10, and S11). The governing equations for these control structures are embedded directly into the system solution matrix, facilitating the use of classical state estimation and control algorithms for LTV systems.

Mobile computational elements: The original numerical scheme suffers from instability when modeling backwater effects. We correct this problem by introducing mobile computational elements that track discontinuities in the water surface profile (see Section S12).

Accuracy improvements: High-gradient conditions were found to induce mass balance errors under the original formulation. To reduce these errors, we re-derive the recurrence relations and superlink boundary conditions to eliminate some error-generating assumptions (see Sections S4 and S5).

---

[2]Note that for sufficiently large state-space systems, A_1 should not be inverted directly [46]. Rather, equations 26, 29 and 31 should be solved implicitly for their dependent variables using algorithms for solving linear systems of equations (e.g. Gaussian Elimination).

Adaptive step size control: Implicit solvers typically maintain better stability at large time steps when compared to explicit solvers—especially during near-steady-state conditions. To exploit this strength, we implement an adaptive time-stepping algorithm that allows the solver to use a small timestep during highly transient conditions and a large timestep during weakly transient conditions (see Section S13) [47].

Infiltration/runoff coupling: To enable the use of `pipedream` as a standalone stormwater management model, we implement an infiltration module based on the well-known Green-Ampt formulation [48]. This module is coupled bi-directionally with the hydraulic solver to enable more accurate representation of overland flow and runoff generation (see Section S14).

3.5 Implementation

The `pipedream` toolkit is implemented in the Python programming language, which provides a powerful interpreter environment for interactive use [49]. Acceleration of numerical code is realized using the *numba* just-in-time compiler, which compiles native Python code to machine code that achieves speeds comparable to code written in C or FORTRAN [50]. We accelerate the solution of the system matrix equation by using a banded matrix solver: upon model initialization, the system matrix is automatically permuted into a banded form using the Reverse Cuthill-McKee algorithm [51]. All code and data for this study are available at `github.com/mdbartos/pipedream`.

3.6 Model validation

We assess the state estimation toolkit by applying our methodology to a real-world stormwater network and then evaluating the extent to which the Kalman filter improves the accuracy of interpolated and forecasted system states. First, a real-world catchment is selected, and real-time depth data is collected at four sites. We then construct a `pipedream` model of the catchment and force the model with a real-world storm event. We then use a holdout cross-validation approach to measure the extent to which fusing sensor data at selected sites reduces error at the holdout sites. We also evaluate the ability of the Kalman filter to forecast system states by fusing sensor measurements at

one-hour intervals and quantifying the reduction in error over the remainder of each hour. These two applications—interpolation, and forecasting—represent important use-cases for real-time state estimation with continuous monitoring.
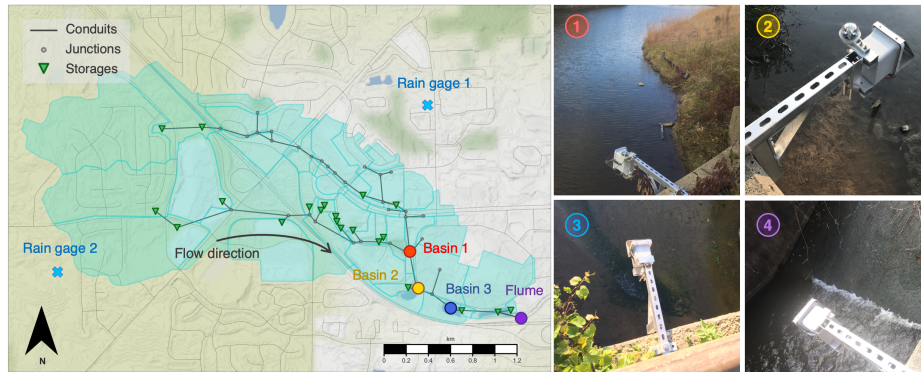

Figure 3: Overview of study area. Left: Map of stormwater network including the contributing area, channel network, hydraulic storage elements, and four sensor sites. Right: Photographs of sensor sites located on basins 1-3 and outlet flume.

Our case study focuses on a 5.85 km^2 urban watershed located in the Midwestern United States (Fig. 3). This watershed is the subject of a long-term monitoring project led by the authors and thus features roughly two years of continuous sensor data [52]. Roughly 47% of the catchment is impervious, with a majority of the impervious area located towards the downstream half of the catchment. The representation of the channel network is derived from survey data and engineering drawings that describe the network topology, hydraulic geometries, storage curves, and various hydraulic and hydrologic parameters needed to properly model catchment dynamics.

To characterize the response of the catchment, wireless ultrasonic depth sensors are installed at four locations in series along the mainstem of the watershed (numbered in increasing order from 1 to 4 in the downstream direction). These sensors continuously report the distance to the water surface at an adaptive sampling rate that ranges from roughly 2 minutes to 1 hour, with a manufacturer-specified reading-to-reading error of approximately 1mm. Sensors at sites 1 and 2 measure the water level in two relatively large retention basins. The sensor at site 3 measures the water level in an outlet box

16

directly downstream of a third retention basin. Overflow from this retention basin drains into the outlet box through a rectangular weir that is approximately 5m wide. Finally, the sensor at site 4 measures the water level at the downstream end of a rectangular flume. The rectangular flume is roughly 5m wide and is connected at the upstream end to a fourth retention basin. Raw sensor data is preprocessed using a combination of manual and automatic filtering techniques. Invalid readings and sensor spikes are flagged and removed using an automated filtering routine consisting of range checks and threshold tests on the second derivative of the sensor signal. After the initial quality control, distance measurements are converted to water depth estimates using field measurements of the sensor offset to the channel bottom.

We apply the new hydraulic model to a real-world storm event occurring on August 8, 2017. This rain event is selected because (i) all sensor sites were active and reporting during this storm, and (ii) the peak depth of the generated hydrograph is close to the median peak depth for the period of available data, meaning that the storm event is representative of a typical storm event in this location. Precipitation intensity data are collected from two weather stations operated by Weather Underground near the catchment of interest [53]. Runoff is generated from the rainfall using the Green-Ampt formulation and then fed into the hydraulic model. For the purposes of this analysis, we assume uniform rainfall intensity over the catchment and use the average intensity between the two gages as input to the model.

4   Results

4.1   Interpolation

Using a holdout cross-validation approach, we find that the Kalman filter is effective at interpolating system states at ungaged locations. Figure 4 shows the result of the holdout cross-validation assessment. For this experiment, the filter is applied to sensor sites 1 and 3, and the output of the updated model is compared with sensor measurements at sites 2 and 4. From Table 1, it can be seen that the filter reduces error at both holdout sites. While the model performs well on its own, the Kalman filter reduces the mean squared error (MSE) at site 2 by 25.5% and at site 4 by 17.9%. For site 2, a majority of the error is reduced at the peak of the hydrograph, while at site 4, a majority of the

error is reduced at the falling limb. Because the filter improves model accuracy even at locations where it is not directly applied, the holdout assessment suggests that the Kalman filter pushes the system closer to its actual state rather than simply "overfitting" individual sites to measured data.

The Kalman filter is capable of correcting error introduced by uncertain dynamical inputs, distinguishing it from a calibration-only approach. From Figure 4, it can be seen that the model without filtering over-predicts the discharge at sites 1 and 2, but under-predicts at sites 3 and 4. This result suggests that the spatial heterogeneity of runoff is a major source of error. This type of error is difficult to counter with model calibration alone, given that calibration tends to target static parameters of the system such as channel roughness coefficients and impervious area percentages. However, spatial variability in runoff may also result from spatial variability in the originating rainfall field. In contrast to continuous calibration, the Kalman filter handles this contingency by correcting system states in real-time, adding and removing mass from the system to match field observations. This approach is robust to both parameter and input uncertainty, making it especially suitable for real-time applications in which the driving input is often uncertain or unknown.
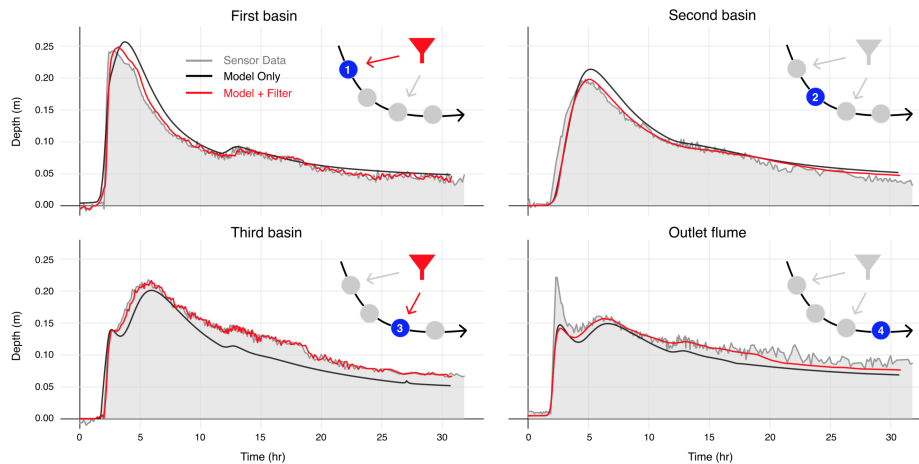

Figure 4: Validation of Kalman filter using holdout assessment. Left: Depth hydrographs at basins 1 and 3, where the Kalman filter is applied. Right: Depth hydrographs at holdout sites where Kalman filter was not applied. The Kalman filter reduces error at both holdout sites.

18

	Basin 2	Outlet Flume
MSE (Model)	3.22E-4	6.36E-4
MSE (Filter)	2.40E-4	5.22E-4

Table 1: Mean squared error of reconstruction at holdout sites.

4.2 Forecasting

By adaptively correcting system dynamics, Kalman filtering pushes the hydraulic system closer to its measured state and thus enables improved forecasting of system behavior. However, not all sites offer the same forecasting benefit, and sensor sites must be selected judiciously to maximize the accuracy of the forecast. Figure 5 shows the result of using the Kalman filter to forecast system states at 1-hour intervals. In this application, sensor data for each site is fused at the first minute of each hour (indicated by circular markers), and the model is then propagated forward in time to forecast system behavior for the remainder of the hour. The effectiveness of the forecast is then gaged based on the reduction in MSE, and the length of time that the correction persists.
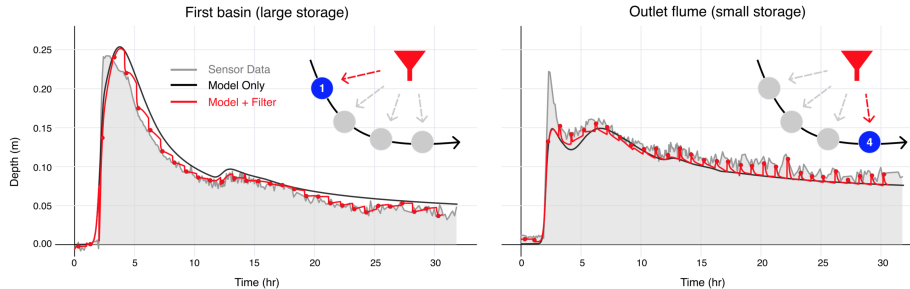

Figure 5: Forecasting hydraulic states at 1-hour intervals using the Kalman filter. Forecasting benefit is greater for sites with large storage capacity (left), and less so for sites with small storage capacity (right).

In general, the most effective sites for forecasting are those for which the volume of storage is large compared to the volume of water entering and exiting the control volume. When the filter is applied to correct the hydraulic head at a large retention pond (site 1), as shown in Figure 5 (left), the effect of the correction persists for the

19

remainder of the hourly interval—in other words, applying the correction significantly changes the trajectory of the hydraulic head compared to the trajectory produced by the model alone. The correction imposed by the filter increases the accuracy of the forecasted states (measured by MSE) compared to the model-only forecast. Sites with larger storage capacity are most effective because small changes in hydraulic head result in relatively large changes to the mass and energy balance of the system, essentially allowing greater control over the state space.

By contrast, sites with smaller storage capacity offer less forecasting benefit, given that the effect of the correction is quickly overpowered by dynamics originating from elsewhere in the system. At the outlet flume (site 4), for example, the proportion of mass entering and leaving the control volume at each time step is relatively large compared to the mass within the control volume itself. Thus, correcting the system state at this location does not significantly alter the amount of mass or energy in the system, and the effect of the correction persists for only a short time after the filter is applied. As shown in Figure 5 (right), the trajectory of the hydraulic head quickly returns to the model-only trajectory after each application of the filter. Thus, sensor data from sites with little storage capacity must be fused at a rapid frequency (on the order of the time step used by the hydraulic solver) in order to produce noticeable forecasting benefit.

4.3 Computational Performance

The `pipedream` model offers significantly improved model run-times compared to pySWMM—an existing interactive solver for stormwater networks. Model run-time comparisons are conducted on networks α, β, and ϵ from the `pystorms` benchmarking library, with network sizes ranging from 26 to 210 nodes (superjunctions) [54]. When compared against pySWMM, `pipedream` ran between 15-190 times faster, with performance gains increasing for larger numbers of computational elements. For many networks, this multiple order-of-magnitude run-time improvement could mean the difference between offline usage and true real-time execution. When compared with SWMM5 in batch execution, `pipedream` is currently about 2.1-3.4 times slower for the benchmark scenarios chosen. However, `pipedream` was able to achieve more consistently stable results. In particular, for the β network, `pipedream` was able to achieve stable results, while SWMM was not— even when using a small timestep (0.1 s). Table 2 shows a detailed breakdown of model

20

| | | | | Model run time (s) | | |
Scenario	Nodes	Links	Duration (hr)	pipedream	SWMM	pySWMM
α	26	30	12	4.42	1.29	64.9
β	210	210	24	42.4	19.98	8210
ϵ	78	77	960	196.0	64.8	10900

Table 2: Performance of model runs under pipedream, SWMM and pySWMM. To assess the performance of pySWMM in an interactive context, pySWMM model runs incorporate the time needed to query depth and flow states from all nodes and links, respectively.

performance for all scenarios. Figures S4-S6 compare model outputs for `pipedream` and SWMM under each scenario.

5   Discussion

By enabling real-time interpolation and forecasting of hydraulic states, `pipedream` provides a powerful new tool for urban flash flood nowcasting. The dynamics of urban flash flooding are complex and spatially heterogeneous, to the effect that there is no existing model that is capable of reliably forecasting flash floods in urban catchments [55]. While many cities use gage networks to help detect flooding and communicate flood alerts, gages are generally restricted to larger streams, leaving significant "blind spots" in the drainage network [56]. To address this problem, the data assimilation methodology presented in this paper will allow emergency managers to better estimate localized flooding at ungaged locations by interpolating hydraulic states from locations where sensor data is available. These high-resolution flood estimates will enable new forms of rapid flood response, such as localized alerts for motorists and targeted dispatch of emergency services. Moreover, by correcting system states in real-time, the `pipedream` toolkit will assist with flood forecasting, enabling emergency managers to more accurately predict downstream flooding at longer lead times—especially in cases where flooding is primarily driven by upstream transport.

In addition to detecting localized floods, the `pipedream` toolkit will also assist in identifying maintenance emergencies. Timely and accurate diagnosis of maintenance issues is essential for effective stormwater management. Short-term maintenance emergencies

(such as storm drain blockages) can lead to localized flooding, while long-term maintenance issues (such as sediment accumulation) can degrade the overall performance of the stormwater system [5]. However, distinguishing true maintenance emergencies from spurious sensor faults remains a persistent challenge—especially when attempting to diagnose anomalies from sensor data alone [57]. `pipedream` handles this ambiguity by combining sensor data with a dynamical model of the stormwater system. By offering two independent estimates of system states, `pipedream` makes it easier to track the source of anomalies and differentiate true emergencies from false positives. Moreover, `pipedream` natively encodes the relative confidence in sensor/model outputs through the measurement/process covariances respectively, allowing users to incorporate prior information about sensor faults directly into the data assimilation procedure.

Perhaps most importantly, the `pipedream` toolkit provides a foundation for real-time control of urban drainage systems. Many cities are now implementing or seeking to implement real-time control systems for urban drainage systems in order to improve performance, cut costs, or mitigate stormwater-related hazards. In both modeling studies and real-world deployments, real-time control has shown proven results in mitigating combined sewer overflows and improving urban water quality [8–16]. However, effective real-time control is predicated on an accurate representation of system states. By fusing sensor data with an accurate physically-based process model, the `pipedream` toolkit provides a firm basis for control, whether for the purposes of simulation or real-world execution. Crucially, `pipedream` integrates control structures into the internal state-space model, facilitating native use of classical control algorithms such as model-predictive control and linear-quadratic regulation.

Conclusions

In this study, we develop a new toolkit for real-time *digital twinning* of urban stormwater systems. This toolkit consists of a robust hydraulic solver based on the full one-dimensional Saint-Venant equations along with an implicit Kalman filtering methodology that facilitates assimilation of real-time sensor data. Drawing on sensor data from a real-world stormwater network, we find that the implicit Kalman filter is effective at both interpolating system states within the network, and forecasting future states based on

22

current measurements. By providing a physically-based methodology for state estimation in stormwater networks, this toolkit will enable system operators to pre-emptively detect and repair blockages, leaks and other maintenance emergencies. Moreover, by improving interpolation and forecasting of system states, our toolkit will provide a strong foundation for model-based real-time control schemes, such as model predictive control and linear quadratic regulation.

Declarations of interest

Declarations of interest: none.

Data availability

Code and data links are available at:

`https://github.com/mdbartos/pipedream`

Permanent DOI: `10.5281/zenodo.3950373`

References

[1] B. Kerkez, C. Gruden, M. Lewis, L. Montestruque, M. Quigley, B. Wong, A. Bedig, R. Kertesz, T. Braun, O. Cadwalader, A. Poresky, C. Pak, Smarter stormwater systems, Environmental Science & Technology 50 (14) (2016) 7267–7273. `doi: 10.1021/acs.est.5b05870`.

[2] E. A. Rosenberg, P. W. Keys, D. B. Booth, D. Hartley, J. Burkey, A. C. Steinemann, D. P. Lettenmaier, Precipitation extremes and the impacts of climate change on

stormwater infrastructure in washington state, Climatic Change 102 (1-2) (2010) 319–349. `doi:10.1007/s10584-010-9847-0`.

[3] B. P. Wong, B. Kerkez, Adaptive measurements of urban runoff quality, Water Resources Research 52 (11) (2016) 8986–9000. `doi:10.1002/2015WR018013`.

[4] G.-T. Blecken, W. F. Hunt, A. M. Al-Rubaei, M. Viklander, W. G. Lord, Stormwater control measure (SCM) maintenance considerations to ensure designed functionality, Urban Water Journal 14 (3) (2015) 278–290. `doi:10.1080/1573062x.2015.1111913`.

[5] J. Wright, D. Marchese, Briefing: Continuous monitoring and adaptive control: the 'smart' storm water management solution, Proceedings of the Institution of Civil Engineers - Smart Infrastructure and Construction 170 (4) (2017) 86–89. `doi:10.1680/jsmic.17.00017`.

[6] R. E. Criss, E. L. Shock, Flood enhancement through flood control, Geology 29 (10) (2001) 875. `doi:10.1130/0091-7613(2001)029<0875:FETFC>2.0.CO;2`.

[7] C. H. Emerson, C. Welty, R. G. Traver, Watershed-scale evaluation of a system of storm water detention basins, Journal of Hydrologic Engineering 10 (3) (2005) 237–242. `doi:10.1061/(asce)1084-0699(2005)10:3(237)`.

[8] L. Montestruque, M. D. Lemmon, Globally coordinated distributed storm water management system, in: Proceedings of the 1st ACM International Workshop on Cyber-Physical Systems for Smart Water Networks - CySWater 2015, ACM Press, 2015. `doi:10.1145/2738935.2738948`.

[9] M. S. Gelormino, N. L. Ricker, Model-predictive control of a combined sewer system, International Journal of Control 59 (3) (1994) 793–816. `doi:10.1080/00207179408923105`.

[10] R. R. van Nooijen, A. Kolechkina, Speed of discrete optimization solvers for real time sewer control, Urban Water Journal 10 (5) (2013) 354–363. `doi:10.1080/1573062x.2013.820330`.

[11] V. Puig, G. Cembrano, J. Romera, J. Quevedo, B. Aznar, G. Ramón, J. Cabot, Predictive optimal control of sewer networks using CORAL tool: application to

riera blanca catchment in barcelona, Water Science and Technology 60 (4) (2009) 869–878. `doi:10.2166/wst.2009.424`.

[12] N. K. Andersen, P. Harremos, S. Sørensen, H. S. Andersen, Monitoring and real time control in a trunk sewer, Water Science and Technology 36 (8-9) (1997) 337–342. `doi:10.2166/wst.1997.0689`.

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

[14] J. R. Middleton, M. E. Barrett, Water quality performance of a batch-type stormwater detention basin, Water Environment Research 80 (2) (2008) 172–178. `doi:10.2175/106143007x220842`.

[15] C. Jacopin, E. Lucas, M. Desbordes, P. Bourgogne, Optimisation of operational management practices for the detention basins, Water Science and Technology 44 (2-3) (2001) 277–285. `doi:10.2166/wst.2001.0780`.

[16] J. F. Carpenter, B. Vallet, G. Pelletier, P. Lessard, P. A. Vanrolleghem, Pollutant removal efficiency of a retrofitted stormwater detention pond, Water Quality Research Journal 49 (2) (2013) 124–134. `doi:10.2166/wqrjc.2013.020`.

[17] G. Karmous-Edwards, Technology perspectives: Digital twin technology becoming a reality for water utilities, Water Finance & Management.
URL `https://waterfm.com/digital-twin-technology-becoming-a-reality-for-water-utilities/`

[18] W. Huber, R. Dickinson, L. Rossman, EPA storm water management model, SWMM5, in: Watershed Models, CRC Press, 2005, pp. 338–359. `doi:10.1201/9781420037432.ch14`.

[19] G. Riaño-Briceño, J. Barreiro-Gomez, A. Ramirez-Jaime, N. Quijano, C. Ocampo-Martinez, MatSWMM – an open-source toolbox for designing real-time control of

urban drainage systems, Environmental Modelling & Software 83 (2016) 143–154. doi:10.1016/j.envsoft.2016.05.009.

[20] pyswmm: Python wrappers for swmm (2017).
URL https://github.com/OpenWaterAnalytics/pyswmm

[21] E. Todini, A bayesian technique for conditioning radar precipitation estimates to rain-gauge measurements, Hydrology and Earth System Sciences 5 (2) (2001) 187–199. doi:10.5194/hess-5-187-2001.

[22] D. Simon, Optimal State Estimation, John Wiley & Sons, Inc., 2006. doi:10.1002/0470045345.

[23] Y. Liu, A. H. Weerts, M. Clark, H.-J. H. Franssen, S. Kumar, H. Moradkhani, D.-J. Seo, D. Schwanenberg, P. Smith, A. I. J. M. van Dijk, N. van Velzen, M. He, H. Lee, S. J. Noh, O. Rakovec, P. Restrepo, Advancing data assimilation in operational hydrologic forecasting: progresses, challenges, and emerging opportunities, Hydrology and Earth System Sciences 16 (10) (2012) 3863–3887. doi:10.5194/hess-16-3863-2012.

[24] D.-J. Seo, V. Koren, N. Cajina, Real-time variational assimilation of hydrologic and hydrometeorological data into operational hydrologic forecasting, Journal of Hydrometeorology 4 (3) (2003) 627–641. doi:10.1175/1525-7541(2003)004<0627:RVAOHA>2.0.CO;2.

[25] G. Thirel, E. Martin, J.-F. Mahfouf, S. Massart, S. Ricci, F. Habets, A past discharges assimilation system for ensemble streamflow forecasts over france – part 1: Description and validation of the assimilation system, Hydrology and Earth System Sciences 14 (8) (2010) 1623–1637. doi:10.5194/hess-14-1623-2010.

[26] A. H. Weerts, G. Y. E. Serafy, S. Hummel, J. Dhondia, H. Gerritsen, Application of generic data assimilation tools (DATools) for flood forecasting purposes, Computers & Geosciences 36 (4) (2010) 453–463. doi:10.1016/j.cageo.2009.07.009.

[27] C. Dechant, H. Moradkhani, Radiance data assimilation for operational snow and streamflow forecasting, Advances in Water Resources 34 (3) (2011) 351–364. doi:10.1016/j.advwatres.2010.12.009.

[28] P. J. van Leeuwen, Particle filtering in geophysical systems, Monthly Weather Review 137 (12) (2009) 4089–4114. `doi:10.1175/2009mwr2835.1`.

[29] V. R. N. Pauwels, G. J. M. D. Lannoy, Ensemble-based assimilation of discharge into rainfall-runoff models: A comparison of approaches to mapping observational information to state space, Water Resources Research 45 (8). `doi:10.1029/2008wr007590`.

[30] D. Schwanenberg, A. van Breukelen, S. Hummel, Data assimilation for supporting optimum control in large-scale river networks, in: 2011 International Conference on Networking, Sensing and Control, IEEE, 2011. `doi:10.1109/icnsc.2011.5874881`.

[31] C. B. Dotto, G. Mannina, M. Kleidorfer, L. Vezzaro, M. Henrichs, D. T. McCarthy, G. Freni, W. Rauch, A. Deletic, Comparison of different uncertainty techniques in urban stormwater quantity and quality modelling, Water Research 46 (8) (2012) 2545–2558. `doi:10.1016/j.watres.2012.02.009`.

[32] B. Schaefli, D. B. Talamba, A. Musy, Quantifying hydrological modeling errors through a mixture of normal distributions, Journal of Hydrology 332 (3-4) (2007) 303–315. `doi:10.1016/j.jhydrol.2006.07.005`.

[33] G. Schoups, J. A. Vrugt, A formal likelihood function for parameter and predictive inference of hydrologic models with correlated, heteroscedastic, and non-gaussian errors, Water Resources Research 46 (10). `doi:10.1029/2009wr008933`.

[34] M. Thyer, B. Renard, D. Kavetski, G. Kuczera, S. W. Franks, S. Srikanthan, Critical evaluation of parameter consistency and predictive uncertainty in hydrological modeling: A case study using bayesian total error analysis, Water Resources Research 45 (12). `doi:10.1029/2008wr006825`.

[35] B. Renard, D. Kavetski, G. Kuczera, M. Thyer, S. W. Franks, Understanding predictive uncertainty in hydrologic modeling: The challenge of identifying input and structural errors, Water Resources Research 46 (5). `doi:10.1029/2009wr008328`.

[36] L. Hansen, M. Borup, A. Møller, P. Mikkelsen, Flow forecasting using deterministic updating of water levels in distributed hydrodynamic urban drainage models, Water 6 (8) (2014) 2195–2211. `doi:10.3390/w6082195`.

[37] M. Borup, Real time updating in distributed urban rainfall runoff modelling.

[38] C. Hutton, Z. Kapelan, L. Vamvakeridou-Lyroudia, D. Savić, Real-time data assimilation in urban rainfall-runoff models, Procedia Engineering 70 (2014) 843–852. `doi:10.1016/j.proeng.2014.02.092`.

[39] A. Breinholt, F. rn Thordarson, J. K. Møller, M. Grum, P. S. Mikkelsen, H. Madsen, Grey-box modelling of flow in sewer systems with state-dependent diffusion, Environmetrics 22 (8) (2011) 946–961. `doi:10.1002/env.1135`.

[40] A. Breinholt, J. K. Møller, H. Madsen, P. S. Mikkelsen, A formal statistical approach to representing uncertainty in rainfall–runoff modelling with focus on residual analysis and probabilistic output evaluation – distinguishing simulation and prediction, Journal of Hydrology 472-473 (2012) 36–52. `doi:10.1016/j.jhydrol.2012.09.014`.

[41] Z. Ji, General Hydrodynamic Model for Sewer/Channel Network Systems, Journal of Hydraulic Engineering 124 (3) (1998) 307–315. `doi:10.1061/(asce)0733-9429(1998)124:3(307)`.

[42] J. A. Zahner, Influence of storm sewers, drainage density, and soil moisture on runoff from an urbanizing catchment, Ph.D. thesis, University of Connecticut (2004).

[43] G. Freni, M. Maglionico, V. D. Federico, State of the art in Urban Drainage Modelling, Tech. Rep. EVK1-CT-2001-00167, University of Bologna (2003).

[44] J. Cunge, Practical aspects of computational river hydraulics, Pitman Publishing Ltd. London,(17 CUN), 1980, 420.

[45] R. E. Kalman, A new approach to linear filtering and prediction problems, Journal of Basic Engineering 82 (1) (1960) 35–45. `doi:10.1115/1.3662552`.

[46] M. Skliar, W. F. Ramirez, Implicit kalman filtering, International Journal of Control 66 (3) (1997) 393–412. `doi:10.1080/002071797224649`.

[47] G. Söderlind, Digital filters in adaptive time-stepping, ACM Transactions on Mathematical Software 29 (1) (2003) 1–26. `doi:10.1145/641876.641877`.

[48] W. H. Green, G. A. Ampt, Studies on soil phyics., The Journal of Agricultural Science 4 (1) (1911) 1–24. `doi:10.1017/s0021859600001441`.

[49] G. van Rossum, Python tutorial, technical report cs-r9526, centrum voor wiskunde en informatica (cwi), amsterdam.

[50] S. K. Lam, A. Pitrou, S. Seibert, Numba: A LLVM-based python JIT compiler, in: Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC - LLVM 2015, ACM Press, 2015. `doi:10.1145/2833157.2833162`.

[51] E. Cuthill, J. McKee, Reducing the bandwidth of sparse symmetric matrices, in: Proceedings of the 1969 24th national conference on -, ACM Press, 1969. `doi:10.1145/800195.805928`.

[52] M. Bartos, B. Wong, B. Kerkez, Open storm: a complete framework for sensing and control of urban watersheds, Environmental Science: Water Research & Technology 4 (3) (2018) 346–358. `doi:10.1039/c7ew00374a`.

[53] Weather Underground, Weather Underground personal weather stations (City of Ann Arbor) (2019).
URL `https://www.wunderground.com`

[54] S. P. Rimer, A. Mullapudi, S. C. Troutman, B. Kerkez, A benchmarking framework for control and optimization of smart stormwater networks, in: Proceedings of the 10th ACM/IEEE International Conference on Cyber-Physical Systems, ACM, 2019. `doi:10.1145/3302509.3313336`.

[55] H. A. P. Hapuarachchi, Q. J. Wang, T. C. Pagano, A review of advances in flash flood forecasting, Hydrol. Process. 25 (18) (2011) 2771–2784. `doi:10.1002/hyp.8040`.

[56] H. Habibi, I. Dasgupta, S. Noh, S. Kim, M. Zink, D.-J. Seo, M. Bartos, B. Kerkez, High-resolution hydrologic forecasting for very large urban areas, Journal of Hydroinformatics 21 (3) (2019) 441–454. `doi:10.2166/hydro.2019.100`.

[57] K. Ni, N. Ramanathan, M. N. H. Chehade, L. Balzano, S. Nair, S. Zahedi, E. Kohler, G. Pottie, M. Hansen, M. Srivastava, Sensor network data fault types, ACM Transactions on Sensor Networks 5 (3) (2009) 1–29. `doi:10.1145/1525856.1525863`.