AGU ADVANCING EARTH AND SPACE SCIENCE

Water Resources Research

RESEARCH ARTICLE
10.1029/2021WR030778

A Reduced Complexity Model With Graph Partitioning for Rapid Hydraulic Assessment of Sewer Networks

Barnaby Dobson[1], Hannah Watson-Hill[1], Samer Muhandes[1,2], Morten Borup[3], and Ana Mijic[1]

[1]Department of Civil and Environmental Engineering, Imperial College London, London, UK, [2]Project Centre Limited–Flood and Water Management Team, London, UK, [3]Krüger A/S, Veolia Water Technologies, Søborg, Denmark

Key Points:

- Automatic graph partitioning can flexibly reduce the complexity of sewer networks to enable surrogate modeling
- CityWat-SemiDistributed can model these reduced networks without needing parameter derivation from high-fidelity simulations
- The combined approach provides computationally cheap simulations and performs accurately even when no high-fidelity model is available

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

B. Dobson,
barnaby.dobson1@gmail.com

Citation:

Dobson, B., Watson-Hill, H., Muhandes, S., Borup, M., & Mijic, A. (2022). A reduced complexity model with graph partitioning for rapid hydraulic assessment of sewer networks. *Water Resources Research*, *58*, e2021WR030778. https://doi.org/10.1029/2021WR030778

Received 14 JUL 2021
Accepted 19 DEC 2021

Author Contributions:

Conceptualization: Barnaby Dobson, Hannah Watson-Hill, Samer Muhandes, Morten Borup, Ana Mijic
Data curation: Barnaby Dobson
Formal analysis: Barnaby Dobson, Hannah Watson-Hill, Morten Borup, Ana Mijic
Funding acquisition: Ana Mijic
Investigation: Barnaby Dobson, Hannah Watson-Hill, Samer Muhandes, Morten Borup, Ana Mijic
Methodology: Barnaby Dobson, Hannah Watson-Hill, Samer Muhandes, Morten Borup, Ana Mijic

Abstract Existing, high-fidelity models for sewer network modeling are accurate but too slow and inflexible for modern applications such as optimization or scenario analysis. Reduced complexity surrogate modeling has been applied in response to this, however, current approaches are expensive to set up and still require high-fidelity simulations to derive parameters. In this study, we compare and develop graph partitioning algorithms to automatically group sections of sewer networks into semi-distributed compartments. These compartments can then be simulated using sewer network information only in the integrated modeling framework, CityWat-SemiDistributed (CWSD), which has been developed for application to sewer network modeling in this study. We find that combining graph partitioning with CWSD can produce accurate simulations 100–1,000× faster than existing high-fidelity modeling. Because we anticipate that many CWSD users will not have high-fidelity models available, we demonstrate that the approach provides reasonable simulations even under significant parametric uncertainty through a sensitivity analysis. We compare multiple graph partitioning techniques enabling users to specify the spatial aggregation of the partitioned network, also enabling them to preserve key locations for simulation. We test the impact of temporal resolution, finding that accurate simulations can be produced with timesteps up to one hour. Our experiments show a log-log relationship between temporal/spatial resolution and simulation time, enabling users to pre-specify the efficiency and accuracy needed for their applications. We expect that the efficiency and flexibility of our approach may facilitate novel applications of sewer network models ranging from continuous simulations for long-term planning to spatially optimizing the placement of network sensors.

Plain Language Summary Existing sewer models are too slow and inflexible to be applied to a range of modern demands that are placed on sewer modeling. For example, climate analyses require simulations of long and continuous timeseries spanning decades over a range of potential futures. The solution to this is to reduce the complexity of sewer models, aiming to provide good enough simulations that are computationally fast and flexible to adapt to the variety of planning situations faced by modern sewer model users. Existing approaches to reduced complexity sewer modeling require time consuming activities such as calibrating to detailed simulations (which are often not available) and inflexible methods to reduce complexity (typically aggregating the network spatially). In this paper, we provide a physically based reduced complexity model that can be applied without calibration. We combine it with graph partitioning, which is a flexible and automated way to reduce a sewer network's complexity in space. Our simulation experiments find the proposed approach to be successful in producing accurate simulations and highlight the utility of reduced complexity modeling by investigating how spatial and temporal resolution interact and impact the accuracy of simulation results.

1. Introduction

Hydraulically based models are used to simulate how sewer networks of urban catchments will respond to precipitation events (Salvadore et al., 2015). These models enable planners to design interventions that might resolve current network issues, and to plan for changes in the urban catchment. However, the models also require the network and catchment to be represented at a high temporal resolution with a fixed spatial representation, resulting in accurate simulations but lengthy simulation times and a lack of flexibility in modeling options. As regulations, such as the UK's newly introduced "Drainage and Wastewater Management Planning," require more computationally expensive applications of sewer network models, such as optimization, real-time control, or scenario analysis to explore the impacts of, for example, changes in climate and land cover (Water, 2019), it is increasingly clear that alternative and more flexible approaches are needed to complement traditional high-fidelity sewer network modeling.

Project Administration: Ana Mijic
Resources: Barnaby Dobson, Samer Muhandes
Software: Barnaby Dobson, Hannah Watson-Hill
Supervision: Morten Borup, Ana Mijic
Validation: Barnaby Dobson, Samer Muhandes
Visualization: Barnaby Dobson, Morten Borup
Writing – original draft: Barnaby Dobson, Samer Muhandes, Morten Borup, Ana Mijic
Writing – review & editing: Morten Borup

Surrogate modeling aims to reduce simulation model run times for computationally expensive applications (Razavi et al., 2012). The principles of surrogate modeling are to apply two broad families of approaches that improve computational efficiency but still enable sufficiently accurate approximations of the model output of interest (e.g., performance metrics). The first is a response surface surrogate, which uses statistical/machine learning techniques to create a data-driven relationship between inputs and outputs, for example, by complementing evaluations of a water distribution network model with a trained neural network to aid optimization (Andrade et al., 2016). Second is low-fidelity modeling, which creates a physically based representation at a reduced complexity, for example, by aggregating in space as a lumped hydrological model would do (Beven, 2012). In this paper, we refer to low-fidelity modeling as "reduced complexity" modeling to avoid the presumption that fidelity may either be just "high" or "low" (and nowhere in between) and to remove the implication that low-fidelity is equivalent to low-trustworthiness.

Following the modern planning demands placed on sewer network analysis, surrogate modeling concepts have recently been applied to sewer network modeling (Li & Willems, 2020; Thrysøe et al., 2019). These existing studies implement a hybrid approach that uses a data-driven analysis of high-fidelity model simulations to derive the parameters of a reduced complexity model. We argue that the requirement to derive parameters using simulations from high-fidelity models limits the applicability of a surrogate approach for sewer network modeling. This is primarily because planning applications of surrogate modeling would typically investigate changes to the physical system, which would then require rerunning the high-fidelity model to derive a new set of parameters. Repeatedly rerunning a high-fidelity model is too labor and computationally expensive a task to perform in a planning context. Thus, we propose that a physically based surrogate approach that does not require derivation of parameters from high-fidelity model simulations is needed. In addition, this could be applied in cases where no high-fidelity model is available at all, and only network information is available, saving a sewer network owner from undertaking the expensive effort of creating and calibrating a high-fidelity model if their need for modeling does not require a highly accurate and detailed representation. Therefore, in this study we perform surrogate modeling of sewer networks using an entirely physically based reduced complexity model to improve usefulness and widen applicability.

Any physically based or hybrid approach for surrogate modeling of sewer networks will require a means to reduce complexity, as mentioned, typically achieved by spatial aggregation. Thrysøe et al. (2019) achieve this aggregation by manual delineation of sub-areas, referred to as "compartments" of a sewer network. Although this allows a user to have some degree of control over the complexity of the surrogate model ultimately created, it is a labor intensive task and will produce different results depending on the person who is performing it. Instead, Li and Willems (2020) extract sub-graphs to identify each self-contained area of a network and lump at this scale. Although this approach is automated and objective, it produces a model of a single, coarse, resolution. While not applied for surrogate modeling, Liu et al. (2018) and Sela Perelman et al. (2015) demonstrate that graph partitioning algorithms can be used to automatically delineate water distribution networks into reduced graphs with a pre-specified complexity. We expect that these graph partitioning techniques might reasonably be used to aggregate sewer networks for modeling, rather than just for administrative purposes [as was done in Liu et al. (2018)]. The combination of partitioning and reduced complexity modeling would give a user the ability to control the spatial complexity of their network model, constituting a far more flexible modeling approach than would be possible with a fixed spatial representation high-fidelity model. In cases where a decision maker is focused on a specific section of network, or where observation data exists that can be used for validation, this approach would enable preserving of locations in the surrogate model that are directly comparable with the real locations.

We also expect that reducing complexity via temporal resolution (i.e., by increasing timestep size) may also be an effective means to perform surrogate modeling of sewer networks. High-fidelity sewer network models are typically not well suited to low temporal resolution due to numerical instabilities that are introduced (Falter et al., 2013). However, surrogate modeling presents an opportunity to formulate simplified, but more accommodating, numerical schemes. Ultimately this will facilitate larger timesteps than would be possible with high-fidelity models, enabling investigating the interactions between spatial and temporal resolution and giving a reduced complexity model user more control of their model complexity to better suit their needs (e.g., for scenario analysis or optimization).

Thus, we propose that, by combining graph partitioning with physically based reduced complexity modeling, we can create simplified sewer network models of a pre-specified efficiency and accuracy that can flexibly

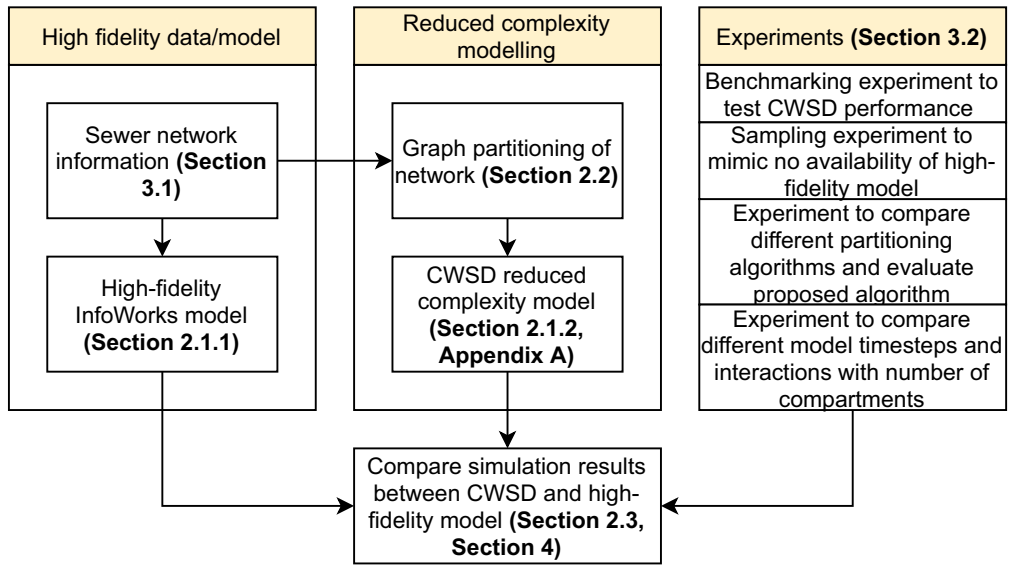

Figure 1. Illustration of the workflow presented in this study.

be adapted to a modeler's requirements. Since none exist, we first present a novel physically based surrogate model for processing flows in an aggregated sewer system that only requires asset information without using high-fidelity model simulations to derive parameters. We also present an experiment to mimic a case where no high-fidelity model is available at all, instead relying on a sensitivity analysis to provide physical parameters that are typically found in a high-fidelity model, but not in raw sewer network data. We then compare a range of graph partitioning algorithms, including two novel approaches tailored to sewer networks, to enable a pre-defined spatial aggregation for surrogate modeling. Finally, we present experiments that investigate how spatial and temporal resolution of surrogate models interacts with both simulation efficiency and accuracy, information that is vital to know in advance of surrogate model applications such as optimization and scenario analysis. Ultimately, we envisage such an approach to be of great value to those that require rapid evaluation of sewer network systems, while also providing further insight into the recurring questions of spatial/temporal resolution in the context of environmental modeling.

2. Methods

To facilitate simulations with a computational efficiency and complexity that suits the needs of a given user, a modeling approach that can accommodate spatial and temporal aggregation is needed. The high-fidelity model against which simulations are compared, InfoWorks, is presented in Section 2.1.1. The novel, reduced complexity model that can make these accommodations, CityWat-SemiDistributed (CWSD), is presented in Section 2.1.2 and Appendix A. The graph partitioning algorithms that enable a user to specify the degree of spatial resolution in CWSD are presented in Section 2.2. We detail the evaluation of CWSD in Section 2.3. The case study catchment, Cranbrook, UK, is described in Section 3.1. The specific experiments that bring together CWSD and partitioning are detailed in Section 3.2. The workflow used in this study is summarized in Figure 1.

2.1. Sewer Modeling at Different Levels of Complexity

2.1.1. High-Fidelity Modeling

A range of high-fidelity models exist for sewer network simulations, for example, InfoWorks ICM (Innovyze, 2014), SWMM (Rossman, 2010), or MIKE URBAN (DHI, 2014). In this study, we use InfoWorks ICM since it is widely used in the UK and in research applications (Babovic & Mijic, 2019; Gong et al., 2018; Muhandes et al., 2021; Ochoa-Rodriguez et al., 2015). In InfoWorks ICM, a conduit is represented as a defined length link in the network between two nodes. The boundary condition between the link and a node is either of the outfall or head loss type. The gradient of a conduit is defined by invert levels at each end of the link. The conveyance

function in InfoWorks can use Manning's equation or the Colebrook-White equation. We select Manning's equation that uses the hydraulic gradient resulted by the upstream and downstream head-difference when the network is surcharged or the conduit slope otherwise. InfoWorks ICM has its own 1D St Venant solvers to undertake hydraulic calculations of a drainage network for transient flow (Innovyze, 2014). For hydrology calculations, InfoWorks ICM uses SWMM5 hydrology and infiltration modules (Rossman, 2010).

2.1.2. CityWat-SemiDistributed (CWSD) Reduced Complexity Modeling

CWSD is an open source integrated urban water modeling framework and software designed for easy setup, efficient computation times achieved by spatial aggregation (hence semi-distributed) and flexibility to accommodate a wide range of system types (Dobson, Jovanovic, et al., 2021). The Python source code used in this study is provided at Dobson, Watson-Hill, et al. (2021). The modeling framework provides a node- and arc-based system representation, where nodes are various elements of the water system (e.g., treatment plants or aggregations of pipe network), and arcs are the means that flow information is passed between nodes (e.g., specific pipes or rivers). The semi-distributed nature of CWSD and its ability to flexibly change how specific subsystems are represented makes it a uniquely suitable tool to explore reduced complexity modeling in a sewer network context. In Figure 2, we demonstrate how a sewer network can be split into "compartments" to provide a reduced complexity semi-distributed representation. As in Thrysøe et al. (2019), the volume of water storage available in specific manholes and pipes can be aggregated to provide a compartment capacity. Thus, each CWSD node in this study is a compartment of semi-distributed sewer network, while arcs between compartments are conceptually identical to the specific pipes in the high-fidelity model that cross the border between two compartments.

In this study, flow information is passed entirely in the form of a push request, that is, a request to push water from an upstream to a downstream compartment. This form of information passing is a common integrated modeling technique in the context of urban wastewater modeling, for example, as in CityDrain3 (Burger et al., 2016). Thus, the simulation takes place by orchestrating a series of push requests between compartments, as demonstrated in Appendix A, Figure A1.

Previously, CWSD has been demonstrated only for extremely low-resolution urban water management modeling (i.e., London's sewer network was aggregated in semi-distributed areas of around 200 km^2). Thus, in this study, we have developed a range of functionalities to enable CWSD for the finer resolutions required for sewer network simulation, namely, pipe travel times, variable timestep size, simple hydraulics, and time-area runoff calculations. This more physically based representation, and the other aggregated compartment-scale parameters needed to implement them, is detailed in Appendix A.1. We summarize the key hydraulic assumptions below:

1. Flow between compartments is governed by Manning's equation. Where elevation head is based on the elevation of sending/receiving manholes in a compartment, and pressure head is determined by the current storage and total chamber area of manholes in a compartment. The area and hydraulic radius terms in Manning's equation are assumed constant at full-bore. This will result in a slight underestimation of travel times during less intense events, but we assess that the negative impact of this is negligible
2. Travel time within a compartment and between compartments is tracked but assumed to travel at the full-bore velocity of pipes. This will result in slight underestimation of travel times during less intense events
3. If push requests from previous timesteps arrive at a flooded compartment, they are rejected and immediately returned to the upstream compartment. This provides a representation of the backwater process but does not account for hydraulic head generated by flooded water
4. Surface runoff is represented using the time-area method as described in Butler & Davies (2004)

We also provide a worked spreadsheet example performing the compartment aggregation process and hydraulic calculations in Supporting Information S1 and S2. To help users understand whether a CWSD approach may be suitable for their case study, we also provide the network parameters that are needed to derive the reduced complexity model in Table 1. We note that pipe roughness and runoff coefficients are not commonly available, and the impact of uncertainty in these parameters is later tested in a sensitivity analysis, described in Section 2.3.

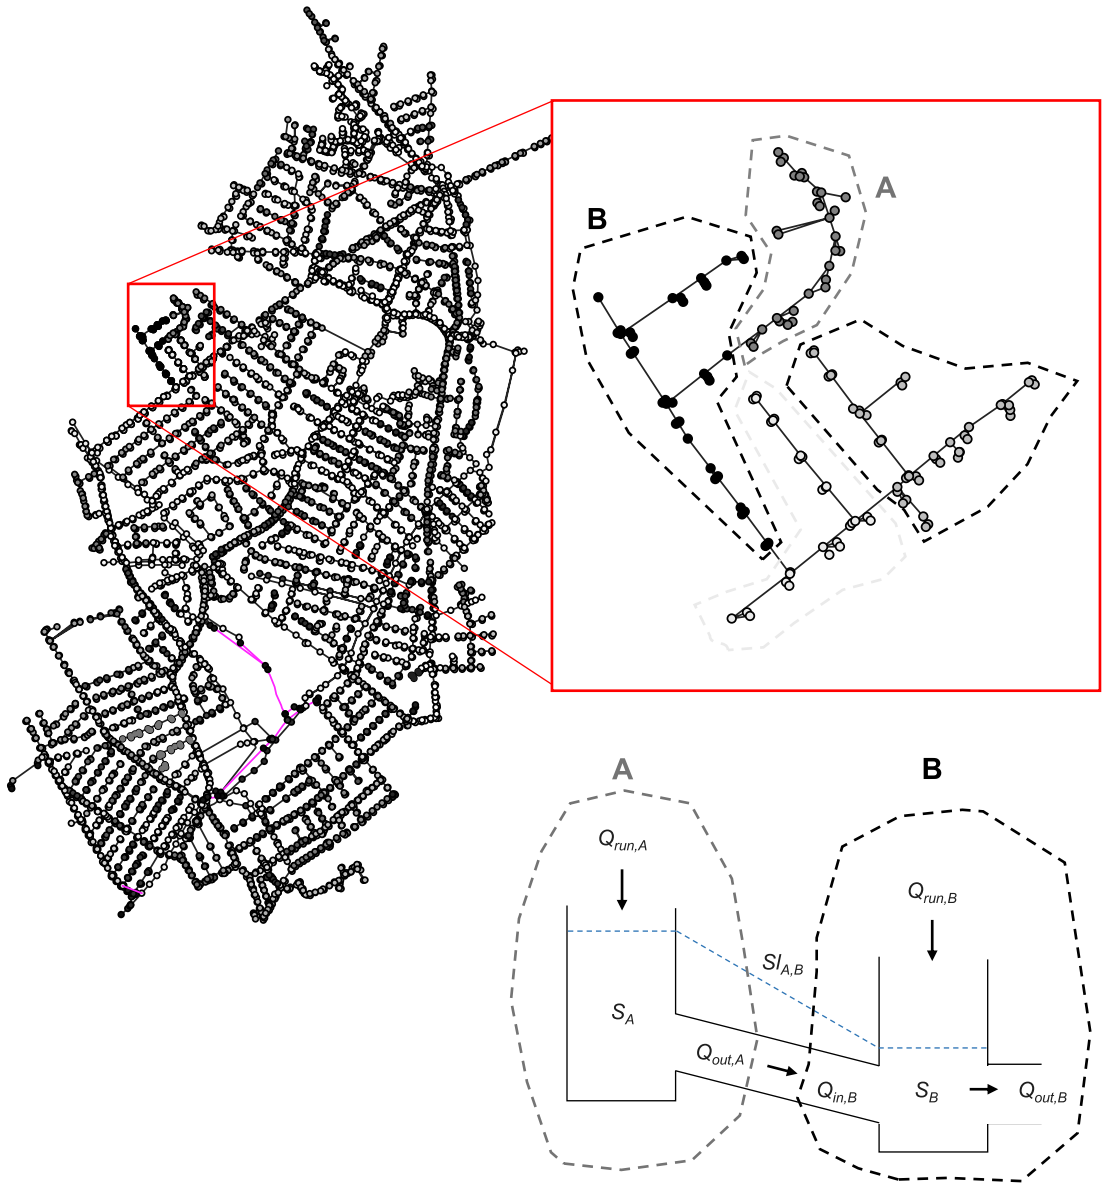

Figure 2. (Map) An example depicting how the Cranbrook sewer network (Babovic & Mijic, 2019) for our case study can be compartmentalized into a semi-distributed, reduced complexity, representation. (Schematic) Q indicates a flow (m^3/s), S indicates a storage (m^3), and Sl indicates a hydraulic gradient (m/m) of the two compartments, A and B. Lakes, weirs, pumps, orifices, headwalls, and culverted watercourses are highlighted in pink.

2.2. Sewer Network Aggregation

2.2.1. Application of Graph Partitioning

As anticipated in the introduction, to perform reduced complexity modeling that can be scaled to large sewer networks and improve the usability of these approaches, an automated method to delineate compartments, spatially aggregating them, is necessary. A key innovation in this study is the application of mathematical graph partitioning techniques to perform this task. The goal of graph partitioning is to allocate sets of nodes into sub-groups, resulting in a graph with a smaller number of edges but retaining the overall connectivity of the original network (Buluç et al., 2016). Because there is a wide range of different available techniques to perform graph partitioning, we compare a selection from different method families: spatial, spectral, and heuristic methods. It is also common to combine multiple algorithms to improve the efficiency of partitioning, called a "multi-level" algorithm (Karypis & Kumar, 1997). We test a selection of these multi-level approaches. We propose two novel multi-level approaches that harness the power of one of the most popular algorithms, Louvain (Blondel et al., 2008), but

Table 1
A Summary of the Asset Data Required to Formulate the CWSD Reduced Complexity Model

Element type	Parameter	Unit
Pipes	Length	m
	Cross section	m^2
	Upstream node	-
	Downstream node	-
	Gradient	m/m
	(Manning's *n* roughness)	-
Sub-catchments	(Runoff coefficient)	%/100
	Area	m^2
	Drainage node	-
Manholes	Chamber area	m^2
	Chamber flood level	m
	Chamber floor level	m
	Invert levels	m
Weirs	Crest height	m
	Width	m
	Discharge coefficient	-
Orifices	Invert level	m
	Discharge coefficient	-

Note. Mannings roughness and runoff coefficient are shown in brackets because, as described in the main text, they are not commonly available, and their impact on simulations can be captured using random sampling.

enable a user to specify the number of groups that it produces to facilitate more widespread application for sewer network aggregation. Specifying the number of groups is key to the flexibility and ability to control model complexity for the approach. Because the total number of groups is inversely proportional to the average surface area drained by a given compartment, the number of partitions can be considered an analogue to the spatial resolution.

The first method, Louvain with spectral reduction, runs a first pass of the Louvain method and then relies on spectral clustering to achieve the desired number of groups. The second method, Optimized Louvain parameters with spectral bisection, instead tunes the parameters of the Louvain algorithm and uses spectral clustering to bisect the largest clusters until the desired number of groups is met. We test both methods with CWSD because visual inspection of the partitioned networks showed more sensible results produced by these methods than the other tested methods.

The partitioning methods used in this study, their family, the source for the algorithm implemented, and a summary are shown in Table 2. We include a shorthand in the "Key" column that results are labeled with.

2.2.2. Preservation of Key Sewer Elements

A common requirement that we anticipate for sewer network aggregation techniques is the ability to preserve user defined key locations within the partitioned graph. In this study, we do this to provide common locations for comparison arcs across partitioning methods (Section 2.3) and to ensure hydraulic structures are represented (Section 3.2). The methods described in Table 2 do not make provisions for this. A simple technique to overcome this is presented below that is compatible with all methods in Table 2.

1. Remove the sewer elements that are to be preserved from the network graph. If an arc is to be preserved, then remove both the up- and downstream nodes
2. Sub-graphs are created from the new network graph, resulting in multiple graphs that are separated by the removed nodes
3. Partitioning algorithms are run separately on each sub-graph. If the number of desired groups is provided (as is the case for all methods in Table 2 except for Louvain), then these are distributed in proportion to the number of nodes in each sub-graph. If sub-graphs contain a small number of nodes, they will still be assigned a minimum of one group, which may cause the resulting number of groups to be greater than the desired number

We note that this method could equally be used to pre-specify compartments (i.e., regions or sub-regions of interest) rather than individual elements. All elements in the pre-specified compartment(s) would be processed as above, and then assigned into their pre-specifications rather than as individual elements.

2.3. Model Evaluation

For model evaluation, we compare CWSD model simulations against the high-fidelity InfoWorks results, as it is common practice in sewer network modeling (Thrysøe et al., 2019). If observations were available, one could use the method presented in Section 2.2.2 to preserve elements with sensors and monitoring data, performing flow comparison against real measured data at these locations.

The two types of timeseries data we compare are pipe flows and total compartment storage. The arcs between compartments are conceptually identical to specific pipes in the InfoWorks model and so a direct comparison can be made. Because specific sewer elements (pipes/nodes) are simulated in the InfoWorks model, rather than compartments, we aggregate InfoWorks in-node, in-pipe, and flooded volume simulation data to compartment scale to create a timeseries of total compartment storage (compartment storage hereafter), which can then be compared against CWSD simulations. We examine compartment storage rather than only compartment surface

Table 2
Summary of the Partitioning Methods Used in This Study

Method	Key	Family	Source (Python package/Reference)	Summary description
K-Means clustering	km	Spatial	Scikit-learn (Pedregosa et al., 2011)	Allocate nodes to groups, iterate to minimize average distance between nodes and group center
Agglomerative clustering	ac	Spatial	Scikit-learn (Pedregosa et al., 2011)	Iteratively merge the two closest nodes until desired number of groups is met
Spectral clustering	sc	Spectral	Scikit-learn (Pedregosa et al., 2011; Shi & Malik, 2000)	Allocate nodes to groups, iterate to minimize the eigenvalues of the graph Laplacian of groups
Fluid communities	fluid	Heuristic	networkx (Hagberg et al., 2008; Parés et al., 2018)	Iteratively propagate groups as fluids through the network, a node's membership is based on the density of fluids in that node
Louvain	louv	Heuristic	Scikit-network (Blondel et al., 2008; Bonald et al., 2020)	Iteratively merge two nodes that will improve the modularity of the network until no more merging can improve the modularity
K-means with spectral bisection	km_sc	Multi-level	Based on Yan et al. (2009)	Preaggregate the network using K-means, then bisect the largest groups using spectral until desired number of groups
Agglomerative with spectral bisection	ac_sc	Multi-level	Based on Yan et al. (2009)	As previous, but using agglomerative in the first step
Louvain with spectral reduction	L_scn	Multi-level	Novel	Iterate steps in Louvain partitioning until there are more groups than desired, then use spectral on the already partitioned network to produce the desired number
Optimized Louvain parameters with spectral bisection	L_L_sc	Multi-level	Novel	Iterate Louvain partitioning with a range of tolerance parameters until there are more groups than desired, then bisect the largest groups using spectral until desired number of groups

water flooded volume because, conceptually, surface water flooding in a compartment must exceed capacity at every compartment sewer element to flood rather than an individual one as is the case for all but incredibly severe real-world precipitation events in which every manhole is flooded. We later provide supporting results (Supporting Information S3) to show that maximum compartment storage is closely related to maximum surface water flooded volume in a compartment, thus justifying that it is still information useful to decision makers who want to understand flooding impacts. We use the Nash–Sutcliffe Efficiency (NSE) performance metric to evaluate both flow and storage simulations.

Comparing different partitioning algorithms with different numbers of desired groups is not trivial to do because there are no common sewer elements for comparison and not all sewer elements are equally important to simulate. We provide a comparison from a user-perspective; selecting arcs that must be preserved during partitioning (see

Section 2.2.2) or compartments that contain specific nodes and comparing only these locations. To account for the importance of different sewer elements, we randomly select a sewer element from each twentieth percentile of flow/storage to preserve. This will ensure that we can determine how different algorithms and numbers of desired groups reproduce behavior in both low and high flow/storage sewer elements. We reproduce all experiments 20 times, randomly selecting different sewer elements to preserve for comparison and to account for variability in how difficult specific different sewer elements are to model. In the main text we provide results for the experiments that compare compartment storage or flow in arcs in the largest twentieth percentile, adjudging these to be of most interest to decision makers, and provide the remainder in Supporting Information S4.

Although the approach presented in this study does not use InfoWorks simulations to derive any parameters, it uses parameters from a high-fidelity model that has been calibrated to observed data. The parameters that are calibrated are runoff coefficients and pipe roughness coefficients. Even though our model evaluation does not compare against observed data, and so does not benefit from this use of calibrated parameters, we believe it is important to demonstrate a case where no high-fidelity model would be available at all because we anticipate this will be a common case for CWSD users. To do this, we derive runoff coefficients from land use maps and test roughness parameters within a range of values. We describe the process for selecting runoff coefficients and roughness parameters in greater detail in Supporting Information S6. The outcome is a range of plausible flow and compartment storage timeseries that we evaluate against InfoWorks and CWSD in the case where it is using parameters from the InfoWorks model.

3. Experimental Setup

3.1. Case Study Catchment

To test the proposed CWSD simulation model and network aggregation techniques, we use the Cranbrook stormwater drainage network in London, UK. The network depicted in Figure 2 drains 8.5 km^2 with over 100 km of pipes, before discharging into the Roding River at two outfall locations. It is ideal for this experiment because it is large enough that manually delineating compartments would be undesirable from a user perspective, but small enough that InfoWorks simulations can be run in high resolution for month-long precipitation timeseries in a feasible amount of time (1 or 2 days) for the purposes of model comparison. The network also contains a variety of complex hydraulic features (lakes, weirs, pumps, orifices, headwalls, and culverted watercourses that drain into the stormwater network, indicated in pink in Figure 2) that will test the ability of CWSD to represent more varied systems. We run the experiments (Section 3.2) for two particularly wet month-long storm timeseries, January 2017 and August 2018. These storms were produced from a radar composite, described in (Dobson, Jovanovic, et al., 2021). The Cranbrook stormwater catchment, and its InfoWorks modeled representation, has been presented and validated in a range of studies (Babovic & Mijic, 2019; Muhandes et al., 2021; Ochoa-Rodriguez et al., 2015).

3.2. Experiments

Because the application of the CWSD simulation modeling framework is novel in a sewer network modeling context, we first perform a benchmarking experiment to ensure that it can suitably be applied to test the performance of different partitioning algorithms. To create a benchmark semi-distributed representation of the network, we use the Louvain algorithm, Table 2. We select this algorithm over a manual delineation, as was used in Thrysøe et al. (2019), to ensure the representation is objective. We select Louvain over other available algorithms because it is one of the simplest to apply, its ubiquity makes it commonplace in benchmarking experiments, and, unlike the other methods, it produces a fixed number of groups so cannot be easily manipulated by tuning parameters. We used a 1-min timestep for CWSD in the benchmarking experiment, in line with the InfoWorks timestep.

We use this same setup in the demonstration of the model that mimics a case where no high-fidelity model would be available. We perform 200 simulations in a sensitivity analysis that examines the sensitivity of results to uncertainty in runoff coefficients and pipe roughness coefficients. Parameter ranges are derived from literature, described in further detail in Table S5 in Supporting Information S6. For each simulation, a randomly sampled runoff coefficient for each land cover type is selected, and the ultimate compartment runoff coefficient derived, and a randomly sampled roughness coefficient is assigned for each model arc.

We next compare the various partitioning algorithms. Since all algorithms in Table 2, except for Louvain, require the user to specify the number of groups produced, we test a variety of different specified number of compartments ranging from 1 to 500. As anticipated in Section 2.3, we randomly preserve sewer elements to ensure like-for-like comparison of the algorithms. To ensure the results are not biased by the random choice of sewer elements to preserve, we repeat the experiment 20 times. We used a 1-min timestep for CWSD in the benchmarking experiment, in line with the InfoWorks timestep.

Finally, to understand how complexity in space and time interact, we run experiments that compare how performance changes depending on the desired number of compartments and CWSD timestep size. Our end-goal in this experiment is for a user to be able to specify their desired model accuracy and be able to use the most efficient reduced complexity network aggregation to achieve that accuracy. As well as varying the number of compartments, as in the second experiment, we also vary the timestep size between thirty seconds and one day. We use the partitioning algorithm that is adjudged to perform best from the second experiment.

Because a key benefit of the approach we present is the ability to specify simulation efficiency, and not just accuracy, we also include additional results that report the speed and pre-processing times of the simulations performed in this work in Supporting Information S7.

Because previous surrogate sewer network experiments [e.g., Thrysøe et al. (2019)] have highlighted the importance of ensuring hydraulic structures are not grouped into compartments, we preserve lakes, weirs, and pumps in the partitioning, as described in Section 2.2.2.

To determine whether the modeling approach can be used to simulate individual sewer nodes (i.e., specific manholes), we also include a supplemental selection of results (S4) where these have been preserved in the same manner as hydraulic structures.

Finally, to avoid timestep to timestep variations in the flow without any hydrological significance obscuring the result comparisons, a 5-min rolling average window has been applied along flows in arcs before plotting or calculating performance measures.

4. Results

4.1. Application of Graph Partitioning and CWSD to Sewer Network Modeling

As described in Section 3.2, the first experiment is to verify that the proposed CWSD simulation framework can produce simulation results comparable to the high-fidelity InfoWorks simulations using the Louvain partitioning algorithm for benchmarking purposes. In Figure 3a, we depict all compartments and arcs colored by their Nash-Sutcliffe Efficiency (NSE) value and sized in proportion to their mean flow or compartment storage. It highlights that larger pipes are accurately simulated, while smaller pipes are more variable. In Figures 3b–3d, we also provide a close-up of key pipe flows through the network on a single day in the month-long simulation.

In Figure 4a, we provide the same map as in Figure 3a, but with simulated compartment storage shown in Figures 4b–4d. We see generally that compartment storage performance is lower than arc performance, although still high in many large storage compartments.

Figures 3 and 4a highlight that performance appears to be influenced by compartment or pipe size. Thus, in Figure 5, we compare maximum pipe flow and maximum compartment storage against NSE performance. We see that, for pipes of low flows (<1 m^3/min, shaded red) and compartments of low storages (<100 m^3, shaded red), reasonable performance (NSE > 0.5) is rare. For average pipes (between 1 and 10 m^3/min, shaded blue) and compartments (between 100 and 1000 m^3, shaded blue), reasonable performance is common, but poor performance (NSE < 0) is also present. For large pipes (>15 m^3/min, shaded green) and compartments (>1000 m^3, shaded green), reasonable performance is common, and poor performance is rare.

As anticipated in Section 2.3, compartment storage represents the total volume of water in nodes, pipes, and flooded in a compartment, which is not conceptually identical to solely flooded volume that decision makers are most likely to be interested in. Thus, Figure S2 in Supporting Information S3, we compare maximum compartment storage as simulated by CWSD against maximum flooded volume as simulated by InfoWorks. We see close agreement between the two, justifying the continued investigation of compartment storage in this study.

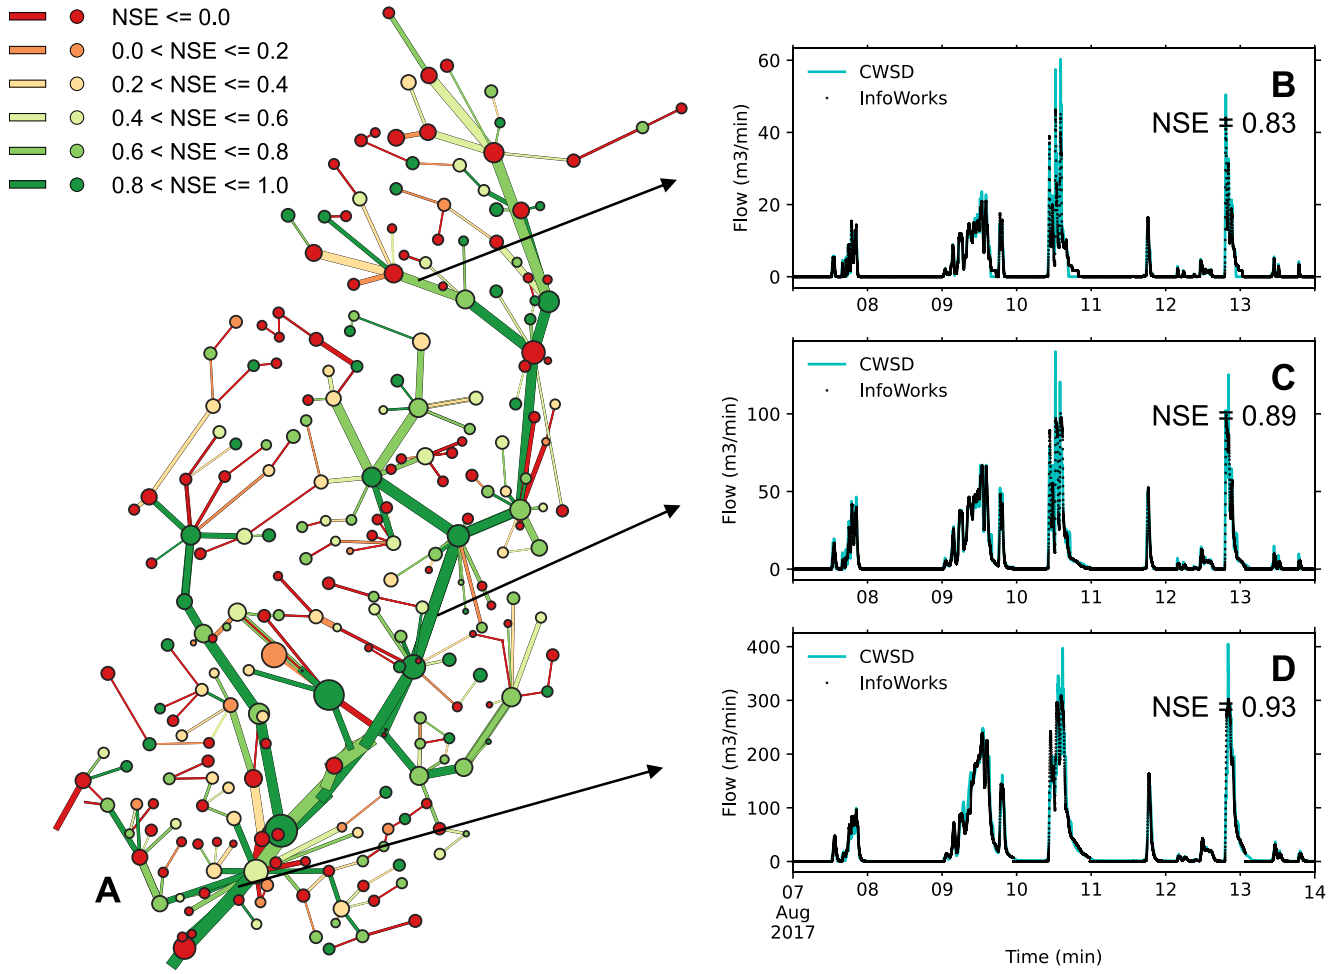

Figure 3. (a) Aggregated network for the benchmark Louvain partitioning algorithm where compartments/arcs are colored by their Nash–Sutcliffe efficiency (NSE) value and sized proportionately to their storage or flow (b)–(d) three sub-sections of flow timeseries along a selection of increasingly large arcs.

4.2. Application of CWSD Without a High-Fidelity Model

In Figure 6, we demonstrate the sensitivity of NSE values to the sensitivity analysis sampling experiment across every arc (Figure 6a) and node (Figure 6b) in the reduced complexity network. This experiment uses raw network data only and tests the sensitivity of results to parameter choices that would otherwise be extracted from a high-fidelity model (runoff coefficients and pipe roughness parameters). In Figure 6a, we see that arcs generally are insensitive to the sampling experiment; if the performance of an arc using network parameters from the calibrated model is high then the performance of that arc is likely to be high using the randomly sampled parameters. There are notable exceptions where the arc performance degrades significantly under the sampling experiment from the performance using the calibrated parameters. These are typically arcs at the edges of the network that have runoff contributed from a single compartment, thus are highly sensitive to changes in runoff coefficients.

In Figure 6b we see that nodes are significantly more sensitive to the sampling experiment, even though the majority of nodes are still insensitive. In general, the larger and more impermeable the sub-catchments within the compartment are, the more sensitive the NSE value is to changing runoff coefficients. In both arcs and nodes, we find that the NSE performance in the calibrated simulations is bounded by the minimum/maximum value from sampling simulations in 60% of arcs/nodes. In 80% of nodes across random sample simulations, the random sample contains at least one simulation with an NSE value greater than the calibrated NSE value, rising to 85% for arcs. In Figure S23 in Supporting Information S6, we provide a set of example timeseries, equivalent to Figures 3 and 4, that highlights how the random sampling experiment generally bounds flow and storage timeseries in addition to the NSE values shown in Figure 6.

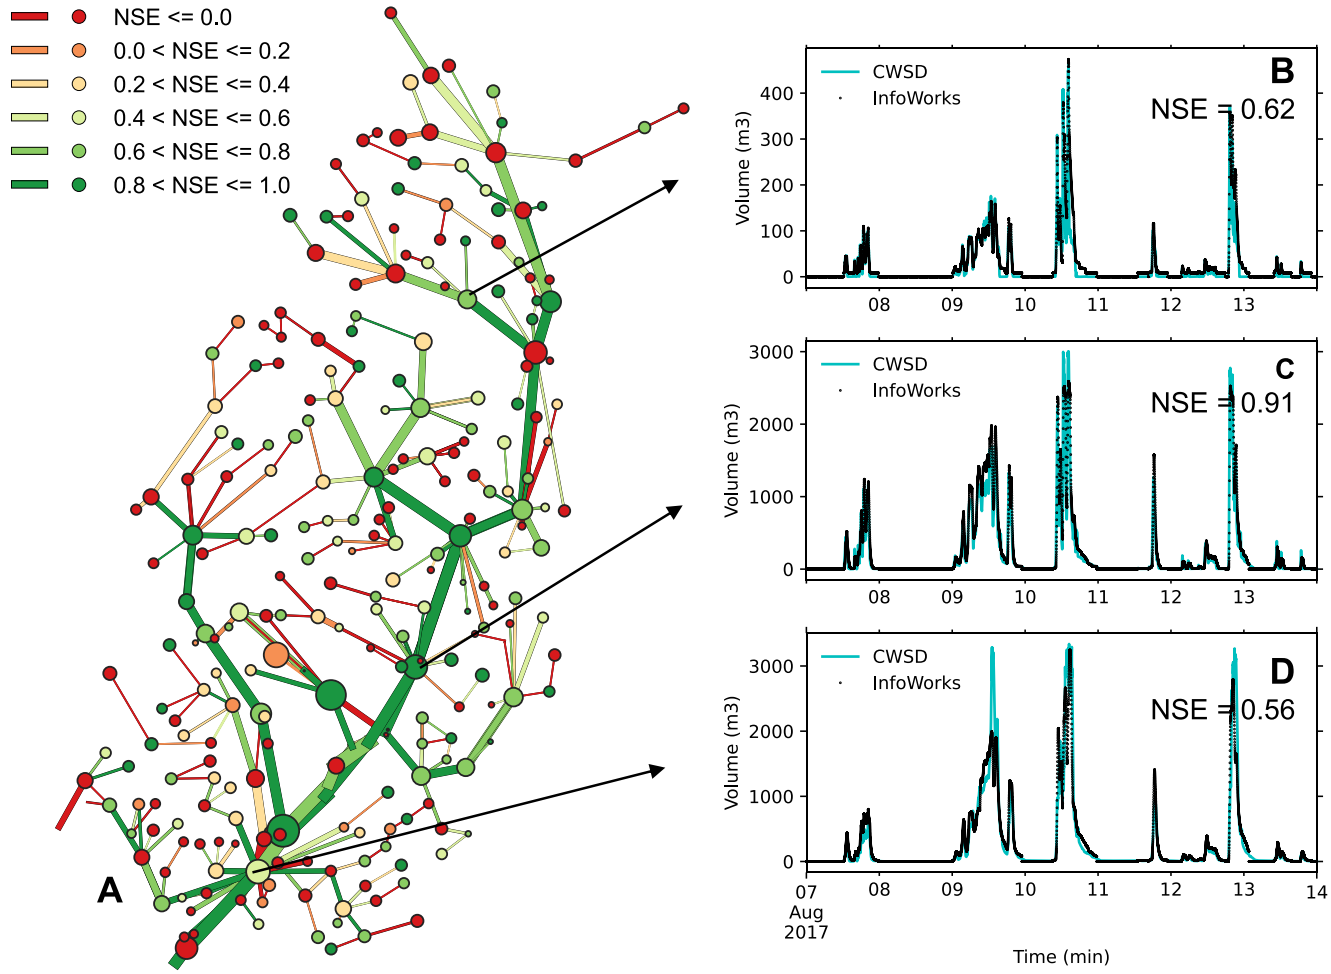

Figure 4. (a) Aggregated network for the benchmark Louvain partitioning algorithm where compartments/arcs are colored by their Nash–Sutcliffe efficiency (NSE) value and sized proportionately to their storage or flow (b)–(d) three timeseries of compartment storage for a selection of compartments.

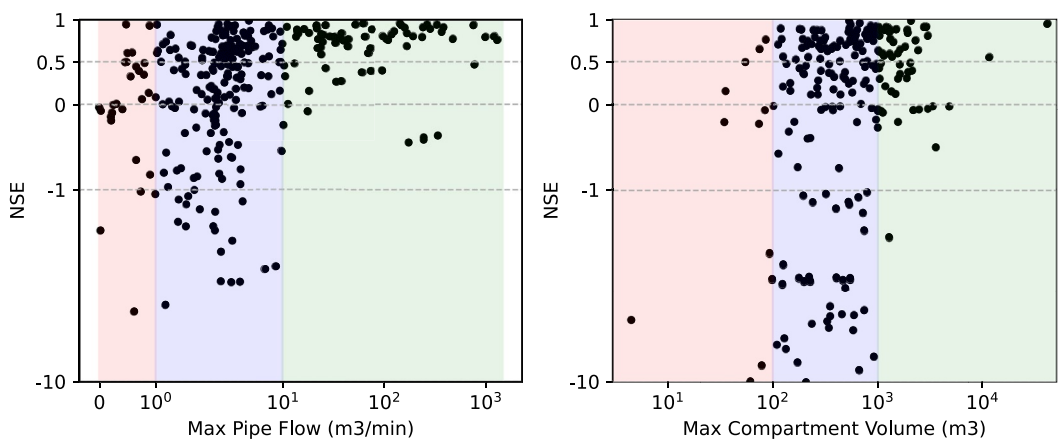

Figure 5. (Left) Nash–Sutcliffe efficiency (NSE) of flow in pipes for the benchmark Louvain partitioning algorithm simulations (*y*-axis) against the maximum flow in that pipe for a given simulation (*x*-axis). (Right) same as (Left) but for NSE of compartment storage. To improve readability, the scale of both axes is linear in the range [−1,1] and log elsewhere.

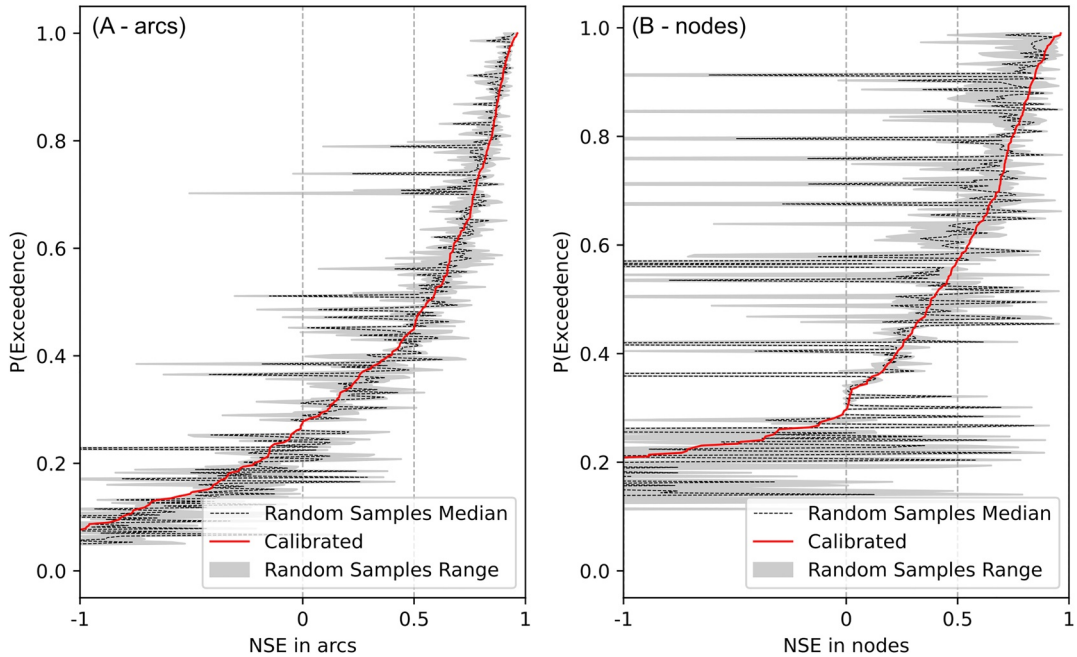

Figure 6. CDF of Nash–Sutcliffe efficiency (NSE) values in (a) arcs and (b) nodes, where NSE is calculated by comparing a CityWat-SemiDistributed (CWSD) simulation against the InfoWorks simulation. The gray band shows the minimum and maximum range of NSE value for a given pipe or node across all random sampling experiments. The dashed black line shows the median NSE value across random samples. The red line shows the NSE of the CWSD simulations using the calibrated parameters from the InfoWorks model. The *x*-axis has been limited to [−1, 1] to improve readability.

4.3. Comparison of Graph Partitioning Methods to Investigate Spatial and Temporal Resolution for Reduced Complexity Models

In this section, we first compare the ability of different partitioning techniques in delineating compartments to provide accurate simulations. As anticipated in Section 2.3, we randomly compare compartments that contain a specific node and flows in preserved arcs in the network to ensure common points of comparison can be made across methods. In Figure 7, we show the NSE values in preserved arcs from the top twentieth percentile pipes (i.e., the pipes with greatest flow). Each panel demonstrates a different partitioning algorithm. Each point is the NSE performance (*y*-axis) of a preserved arc for a given number of specified compartments (*x*-axis). The *x*-value of each point is randomly moved within ±5% of its actual value to better distinguish the density of points, that is, random jitter. Of the compared partitioning methods, Louvain (louv) and the two novel methods, Louvain with spectral reduction (L_scn) and optimized Louvain with spectral bisection (L_L_sc), produce reasonable (>0.5 NSE) median performance metrics (dashed blue lines) when the total number of compartments is greater than 100. All partitioning methods produce inferior (<0.5 NSE) median performance when the total number of compartments is less than 100. We see that some preserved arcs have high NSE values (e.g., the arc highlighted in cyan points) across all simulations, and some have low NSE (e.g., the arc highlighted in red points). The network surrounding the arc highlighted in both cyan and red is shown in Figure S21 in Supporting Information S5. It shows that the cyan arc is a shallow gradient pipe at the end of a linear stretch of network, while the red arc is a steep gradient pipe among a complex tangle of network.

In Figure 8, we show NSE values of compartment storage for compartments that contain a randomly selected sewer network node from the top twentieth percentile nodes (i.e., nodes with the greatest storage). Of the compared partitioning methods, Louvain (louv) and the two novel methods, Louvain with spectral reduction (L_scn) and optimized Louvain with spectral bisection (L_L_sc) produce reasonable (>0.5NSE) median performance metrics (dashed blue lines) when the total number of compartments is less than 200. All partitioning methods produce inferior (<0.5 NSE) results when the total number of compartments is greater than 400. Although NSE performance for compartment storage is, on average, lower than for arc flows, it is below 0 less frequently than for arcs. We see in general some compartment storage timeseries are easier/harder to predict across all methods

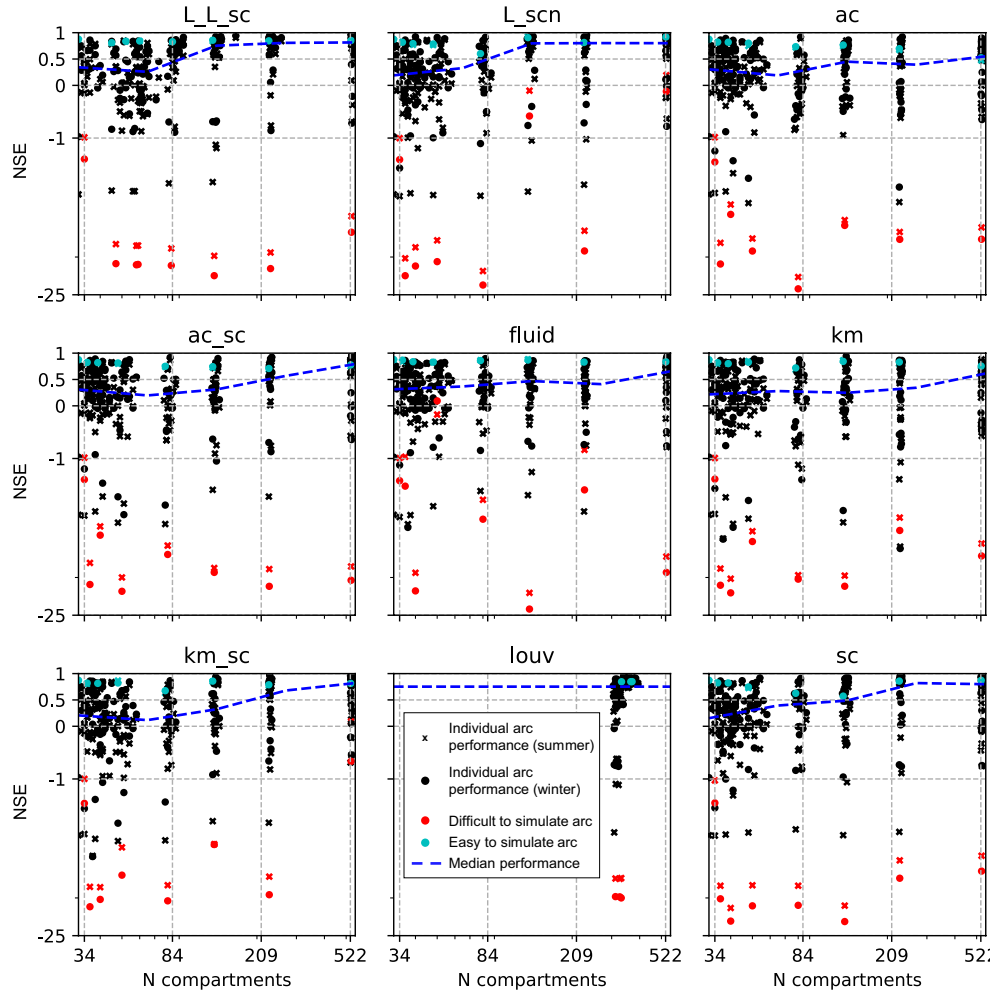

Figure 7. Simulation performance along preserved arcs (points). Interactions between different clustering methods (panels) and different number of compartments (*x*-axis) are shown. 2 month-long precipitation timeseries have been used (cross and circles). The median performance of preserved arcs is plotted (blue dashed lines). The worst performing (red points) and best performing (cyan points) preserved arcs are highlighted. Random jitter in the *x*-axis has been added to better distinguish between points. The *y*-axis has been limited to [−25, 1] to improve readability, and the scale is linear in the range [−1,1] and log elsewhere.

(red points are generally lower while cyan points are generally higher); however, this is not as ubiquitous as with arcs (i.e., we also see some high NSE values for red points and some low NSE values for cyan points). In Figure S22 in Supporting Information S5, we see that the preserved red node is near the culverted watercourses in the catchment, while the area around the preserved cyan node is paved.

Following the results that the proposed "optimized Louvain with spectral bisection" partitioning algorithm provided the best results over a range of different number of compartments, we use it to investigate interactions with timestep size. These interactions are displayed in Figure 9a for NSE values in the random sample of preserved arcs from the top twentieth percentile pipes (i.e., the pipes with greatest flow). We see clearly that performance degrades as both the total number of compartments decreases and as the timestep increases. It is noteworthy that reasonable performance (>0.5 NSE) can be attained with a timestep as large as 1 h. There is a marked difference in performance between total number of compartments less than and greater than 100. In the Supporting Information S4, we demonstrate that performance degrades as pipe size decreases, and there is less evidence of clear interactions between timestep and total number of compartments. In Figure S19 in Supporting Information S4, we also demonstrate that this figure is not sensitive to whether individual sewer elements are preserved for comparison or not.

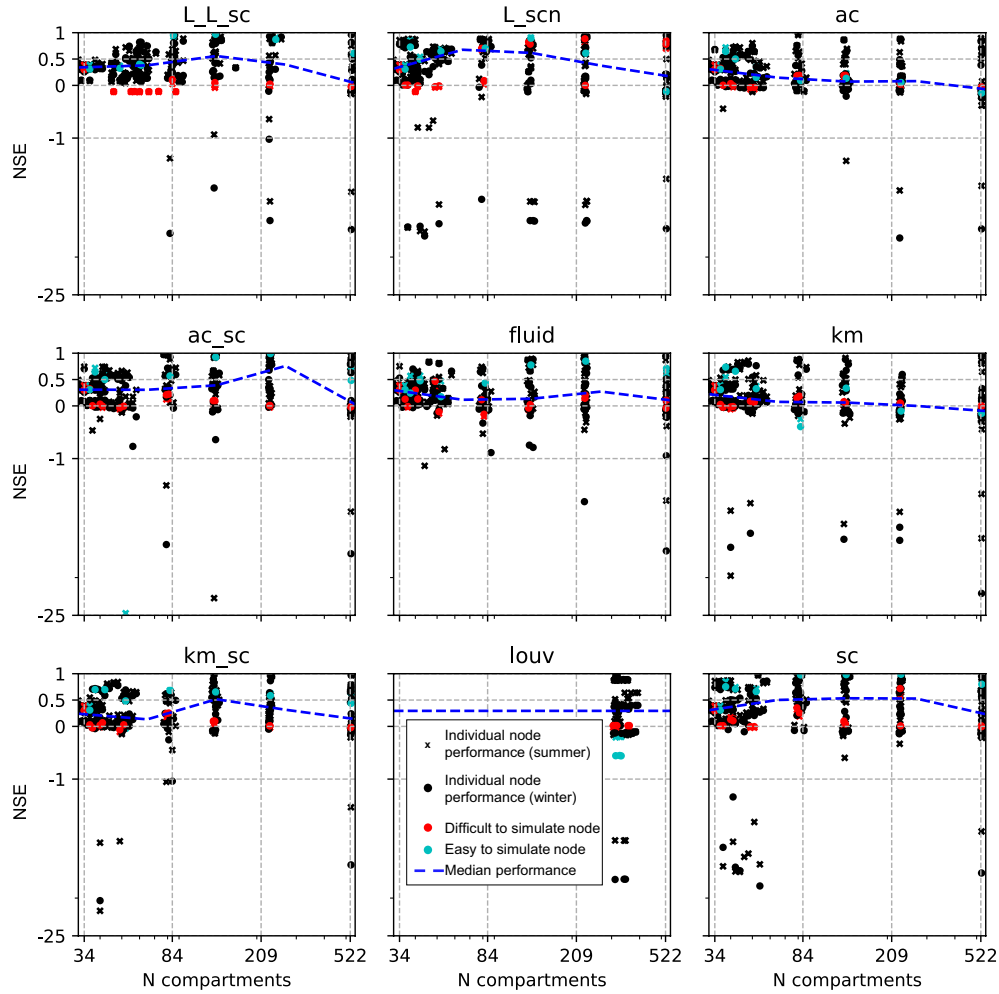

Figure 8. Same as Figure 7, but for compartment storage rather than flow in arcs.

In Figure 9b, we display the NSE values for compartment storage in the random sample of compartments selected from the top twentieth percentile nodes (i.e., the nodes with the greatest storage). Performance degrades quite dramatically once the timestep exceeds one hour, following a similar trend to Figure 9a, albeit with a steeper drop-off in performance. In contrast to Figure 9a, the interaction with total number of compartments seems to prefer around one hundred compartments, with performance degrading as the number deviates from this. We see performance peak at 132 total number of compartments, degrading as the total number increases or decreases. This appears to be due to some aggregation effect provided by compartmentalizing, since it is not seen in the supporting experiment where individual sewer elements are preserved, Figure S20 in Supporting Information S4. In Supporting Information S4, we see that node performance degrades as node size decreases; however, this degradation is to a significantly lesser extent than for the arcs.

5. Discussion

The first aim of this study is to develop a modeling approach to automatically create a reduced complexity simulation model for hydraulic simulations that is suitable for rapid assessment of sewer networks without the need for lengthy calibration by a high-fidelity model. In Figures 3 and 4, we show that graph partitioning can automatically reduce the spatial complexity of a sewer network and that the improvements to CityWat-SemiDistributed made in this study facilitate reasonable simulation results in comparison to a high-fidelity model. We expect that further improvements can be made to the approach. Figures 3 and 4 demonstrate that smaller compartments or pipes closer to the edges of the network typically have worse performance, which we expect is because the behavior in these

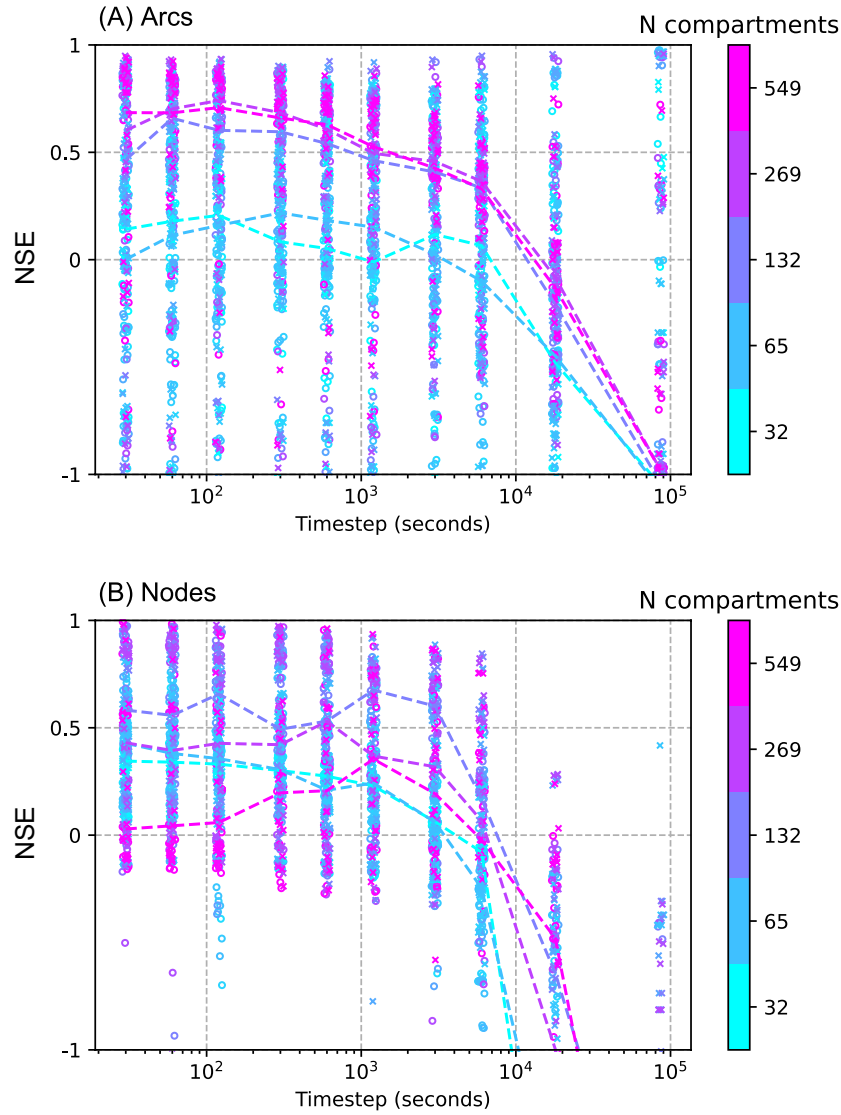

Figure 9. Simulation performance for flows along preserved arcs (points in top panel A) and for compartment storage in compartment nodes (points in bottom panel B). Interactions between total number of clusters (color) and model timestep (x-axis) are shown. 2×1 month-long precipitation timeseries have been used (cross and circles). The median performance of preserved arcs is plotted with dashed lines. Random jitter in the x-axis has been added to better distinguish between points. The y-axis has been limited to $[-1, 1]$ to improve readability.

elements is more sensitive to complex hydraulic processes. Figures 7 and S21 in Supporting Information S5 show that a steep arc in a particularly complex area of network is consistently challenging to simulate. Figures 8 and S22 in Supporting Information S5 show that a compartment that contains culverted watercourses is also difficult to simulate. These findings imply that improvements to the simplified physical representation in CityWat-Sem-iDistributed, presented in Section 2.1.2 and Appendix A, may be needed in areas with more complex hydraulic behaviors, which will be addressed in future work. They also highlight that, while the approach presented can give reasonable performance, it is not suitable for the detailed design of pipes and manholes that a high-fidelity hydraulic representation is required for.

Future work to best demonstrate the full value of the proposed simulation approach would be to derive the reduced complexity CWSD model from network data only and validate by observed data rather than against a high-fidelity simulation model. An open-source case study that could explore the simulated flows at monitored locations could be performed using the data collected for the Bellinge urban water system (Pederson et al., 2021). The demonstration of our proposed approach without using parameters from the high-fidelity model, Figures 6

and S23 in Supporting Information S6, show that simulations are surprisingly insensitive to runoff coefficient and roughness parameter choices. These results are encouraging for those who would apply CWSD with graph partitioning in networks where no high-fidelity model is available. In these cases, we would recommend performing a sampling experiment, as we perform here, to bound the likely true simulation. Identifying which few parameters the model is sensitive may be an opportunity for planners to locate where their interventions will have the most impact.

The next aim of this study is to test a range of graph partitioning algorithms for suitability in reduced complexity sewer network modeling. Figures 7 and 8 show that results are highly dependent on the algorithm selected (dashed blue lines between different panels) and that the novel proposed algorithms (top left and top center panels) in this study perform best. Making broad comments about families of partitioning algorithms is difficult, approaches within the heuristic and multi-level families both contain reasonable and poor algorithms. However, exclusively spatial approaches consistently perform poorly, indicating that a graph partitioning approach must, at a minimum, consider the connectivity of the network.

We demonstrate that a variety of reasonably performing algorithms can be used to pre-specify nodes and arcs in a network to preserve as an isolated element. Although this is useful in our study for preserving hydraulic structures and ensuring that common arcs can be compared across methods, we also expect that this would be useful to those who have key monitoring points or whose goal is to optimally locate monitoring points to maximize "value-for-information." Combined with the dramatic speed gains shown in Figure S24 in Supporting Information S7, we believe the approach presented in this study is ideally suited to a monitoring placement optimization application.

The final aim of this study is to investigate the impact of spatial and temporal resolution on model performance. We first discuss spatial resolution. Figure 5 shows that the model performs best for larger arcs (pipes) and compartments (aggregated areas of sewer network). This indicates, quite intuitively, that sewer elements whose behavior is dependent on a larger area, or a larger collection of other sewer elements, are those whose behavior is best captured by reduced complexity modeling. The seemingly unintuitive result seen in Figure 9b, that compartment storage performance degrades as the total number of compartments increases from 100 to 500, is also explained by this finding. We can assume that the compartment storage performance improves between 30 and 100 compartments because, as shown in Figure 9a, arc performance increases dramatically between 30 and 100 compartments, leveling off after 100. We can verify this by observing Figure S20 in Supporting Information S4, an experiment in which individual manholes were preserved and compared for node storage rather than entire compartments, where the effect is not seen. This implies that reduced complexity modeling requires getting the right balance between capturing spatial heterogeneity (i.e., getting good performance of arcs) and benefitting from the smoothing provided by aggregation (i.e., simulating components that aggregate behavior over larger areas).

From a temporal resolution perspective, Figure 9 demonstrates that the presented approach can return reasonable performance in arcs and at compartments across a range of timestep sizes, between 30 s and 1 h. Salvadore et al., (2015) show that the range of resolutions (both temporal and spatial) covered by this study is not unprecedented for any individual application. However, we cannot identify any sewer network simulation approach that can seamlessly navigate the entire range in a self-contained framework as we do in this study. We believe that this flexibility will be of great utility to users because one may start with a desired simulation time in mind and identify their options for spatial/temporal resolution accordingly, as we show in Figure S24 in Supporting Information S7. We expect that these fast and flexible simulation speeds may facilitate new applications for sewer network simulation models under continuous simulations that span far longer periods than can currently be simulated. Although summarizing the speed-accuracy gains depends on what nodes/arcs/timesteps are under study, Figure 9 demonstrates that a median NSE ≥ 0.5 in both nodes and arcs can be achieved with around 130 compartments and a timestep size of 10^3 s, Figure S24 in Supporting Information S7 shows this to have a simulation speed of around 0.02% that of the InfoWorks ICM simulation speed.

To improve the usability of the methods presented in this paper, we suggest a simple approach toward applying graph partitioning in other case studies. Start with the basic Louvain partitioning algorithm, we have found that this generally gives good enough results and does not require specifying the number of compartments. If there are specific points of interest (e.g., observations or areas where intervention is planned) or hydraulic structures

(e.g., pumps or weirs), these should be preserved as set out in Section 2.2.2. If greater fidelity or computational efficiency is required, then a range of timestep sizes and the proposed L_L_sc partitioning method (Table 2) should be used to experiment with number of compartments, identifying the trade-off between efficiency and fidelity that best suits the user's needs.

6. Conclusions

Modern applications of sewer network modeling, such as optimization and scenario analysis, require far faster simulation speed than can be provided by the existing suite of high-fidelity models that are commonly used by water companies. In this study, we demonstrate the combination of graph partitioning for spatial aggregation and the CityWat-SemiDistributed modeling framework for hydraulic simulations. This calibration free approach can produce accurate simulations, particularly for larger pipes in the sewer network, at a fraction ($<<1\%$) of the computational time of a high-fidelity model and has potential to be used for surrogate model development from network data only with validation against observations. Spatial and temporal resolution in this approach is entirely flexible and can be associated with a deterministic simulation time. We believe that the techniques presented in this study will open a wide range of novel applications for sewer modeling that have previously been prohibited by slow simulations and inflexible modeling approaches.

Appendix A: CityWat-SemiDistributed (CWSD) Sewer Modeling

In this appendix, we first describe the fundamentals of integrated modeling in CWSD and then how the CWSD modeling framework has been applied for sewer network modeling. We then describe the physical processes that have been added and describe how they have been represented for our semi-distributed methodology. The term "sewer junction" is used to represent all the manholes, breaks, lakes, and other node elements that make up the sewer network.

A1. Integrated Modeling With CWSD

As described in Section 2.1.2, CWSD is a node- and arc-based integrated modeling framework and software. Nodes represent physical system components or some semi-distributed aggregation of the system in question (i.e., a compartment), while arcs represent flows between nodes. Nodes and arcs are categorized into a variety of types, each type providing generic functionality for the kind of system component that it represents. The node and arc types used in this study and their functionality are described below:

1. Arcs track both incoming and outgoing water between two nodes. The arc capacity is determined using hydraulic gradient (Section A2), and the travel time is tracked before water can enter a downstream node (Section A3). The travel time along an arc is fixed at the velocity of a pipe's full-bore flow capacity
2. Sewer nodes represent either individual sewer junctions if they have been preserved for comparison purposes (as described in Section 2.3) or to maintain hydraulic structures (Section 3.2), or they represent an aggregation of sewer junctions and pipes within a compartment
3. These nodes have a storage property that tracks how full they are. Their capacity is the total available capacity of all sewer junctions and pipes in the compartment that the node represents
4. The key functionality of these nodes is to receive and discharge water. Received water is tracked to represent travel time across the compartment (Section A3) before being added to the node's "active storage," which is available for discharge downstream. A node's active storage is first discharged to downstream nodes in proportion to available capacity along outgoing links. After this, if the active storage still exceeds node capacity, then the excess is discharged to the adjoining land node as flooded water. A sewer node's total storage (both active and traveling within compartment) and its adjoining land node's flooded volume make up the "total compartment storage" variable used for head calculations (Section A2), and to assess model performance at nodes in Section 4
5. Land nodes represent the sub-catchments that drain into an adjoining sewer node
6. These nodes have a storage property that tracks the flooded volume from spilled sewer nodes or water that cannot be received by sewer nodes
7. Their key functionality is to read precipitation data and apply the time-area method to determine runoff (Section A3)

8. Outfall nodes represent the two outfalls of the catchment. They track received water for mass balance checks

For a given timestep, first land nodes read precipitation and send runoff to sewer nodes. Second, starting from the most upstream sewer nodes and working downstream, sewer nodes discharge water. Finally, at the end of every timestep, mass balance checking is performed, both for each node individually and for the network as a whole. We summarize this process in Figure A1.

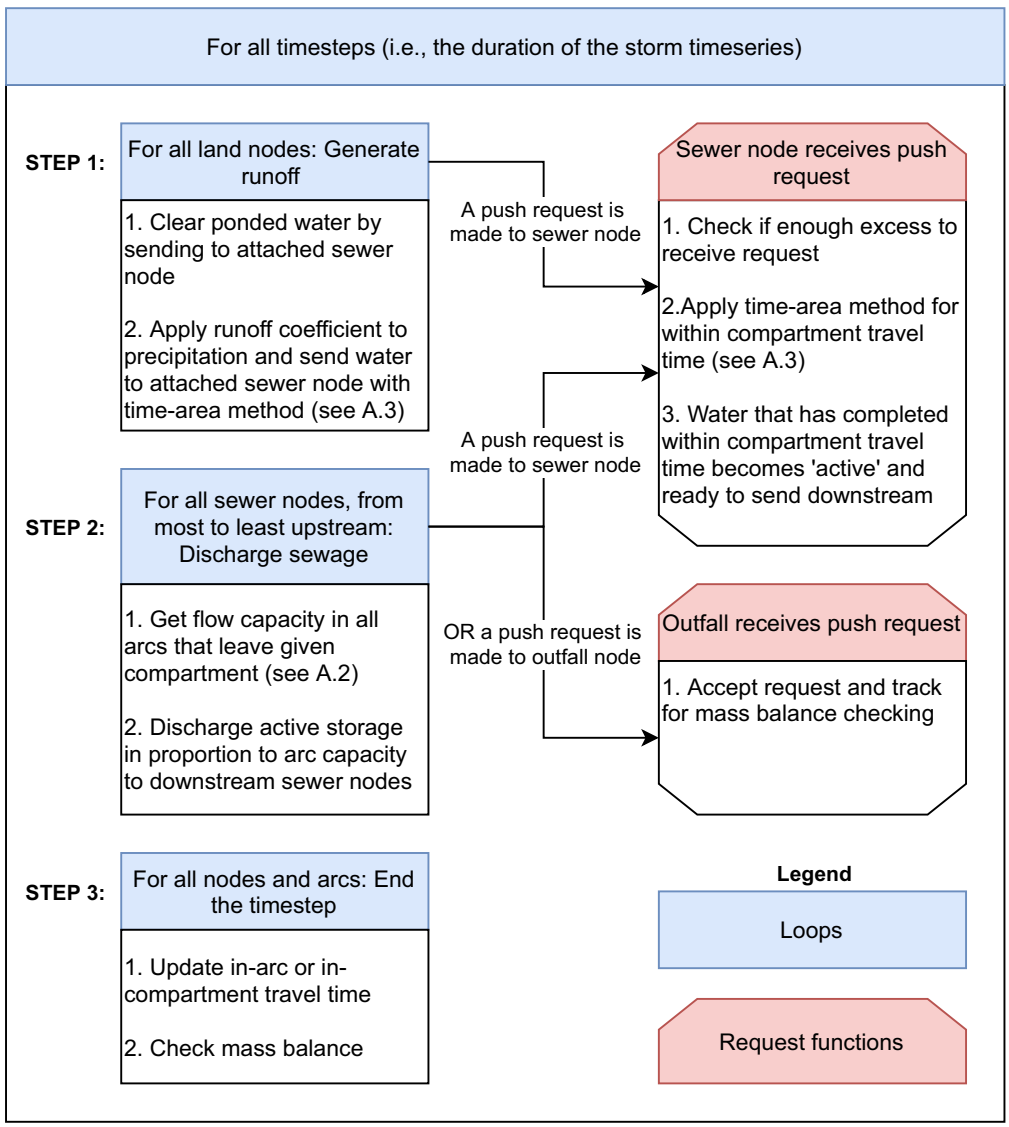

Figure A1. Schematic highlighting the orchestration of nodes during a timestep of simulation in this study's application of CityWat-SemiDistributed (CWSD).

A2. Hydraulics for Between Compartment Pipe Flow

Implementation of pipe hydraulics in this study was reasonably simple because the pipes between compartments are conceptually identical to specific pipes from the high-fidelity model. Thus, the parameters of those pipes (roughness, length, slope, cross-sectional area) remain the same, regardless of how a compartment is aggregated. These parameters enable Manning's equation to be evaluated to derive flow, which determines how much water may travel between two compartments. Manning's equation is given by Equation A1,

$$Q_{i,j} = \frac{A_{p,i,j} R_{i,j}^{\frac{2}{3}} Sl_{i,j}^{\frac{1}{2}}}{n_{i,j}} \tag{A1}$$

where Q is the flow, A_p is pipe cross-sectional area, R is the hydraulic radius, Sl is the slope of the hydraulic gradient, n is the Manning's roughness coefficient along the pipe i,j flowing between compartment i and compartment j. To simplify calculations, A_p and R of pipe i,j are assumed constant at its full-bore flow, thus implying flows are larger than in practice when pipes are not full. Sl is calculated using (Equation A2),

$$Sl_{i,j} = \frac{H_i - H_j}{L_{i,j}} \tag{A2}$$

where H is water head at a compartment and L is the pipe's length. While L is simply the length of the pipe i,j in the high-fidelity model, H is an aggregated value that represents water head at a compartment scale. It is described in Equation A3,

$$H_i = \begin{cases} z_i, & V_i \leq V_{p,i} \\ z_i + \frac{V_i - V_{p,i}}{A_{c,i}}, & V_i > V_{p,i} \end{cases} \tag{A3}$$

where z is the floor elevation of compartment i, V is the volume of water currently in i, V_p is the total volume capacity of pipes in the compartment, and A_c is the total chamber area of sewer junctions in a compartment. Compartments have two possible values for z, an incoming and outgoing value. The incoming z is the average elevation of sewer junctions that receive water from pipes entering the compartment, while the outgoing z is the average elevation of sewer junctions from which pipes leave the compartment. As highlighted in Equation A3, volume of water in a compartment that is less than the total volume of pipes in a compartment is assumed not to generate additional head.

The Cranbrook network used in this study has nine weirs. Among these, there are a variety of weir types and so we do not reproduce each equation separately here, however the same weir equation and parameters as in the high-fidelity modeling is used if an arc between compartments is also a weir. The head calculation for these weir equations is performed the same as in Equation A3.

A worked example that demonstrates these equations for a simplified network is presented in Supporting Information S1 and S2.

A3. Travel Time With Flow Time Information Packaging (FTIP)

Implementing travel time in network models is not trivial in purpose-built software, and this difficulty is exacerbated in an integrated modeling framework where nodes and arcs are discrete from each other, but computational efficiency aims to be preserved. Thus, we introduce a parsimonious approach that is designed to be quick to simulate, easy to understand, flexible to apply and as physically based as possible. We term this technique as "Flow Time Information Packaging" (or FTIP). Using FTIP, every time water is sent somewhere by the model, it is sent in a package (implemented as a dictionary data structure) containing both the amount of water and an integer tracking remaining travel time. Every time a simulation completes a timestep the travel time integer is decremented by one until it reaches zero, at which point the arc attempts to send the FTIP package to the downstream compartment. If a downstream node returns that water could not be received, then that rejected water is assumed to have backed up and returns to the upstream node. Although this accounts for backwater, it prevents flooded sewer nodes from receiving more water. In practice flooded sewer nodes can still receive water if the hydraulic gradient permits however, we considered that calculating the head generated by flooded water above a sewer node would overly increase the complexity of the model, running counter to its ultimate purpose. This will cause underestimates of downstream node flood volumes and overestimates of upstream node flood volumes; however, we consider the effect to be negligible and will not impact estimates of flood occurrence.

Implementing a typical pipe with FTIP requires translating some physical concepts into more discrete terms that can be simulated efficiently. We describe these below:

1. Since FTIP tracks volume and not flow, for pipe capacity calculations FTIP packages must be translated into flow terms. For example, a pipe with a capacity of 10 m^3/s can accommodate two FTIP packages with a volume of 5 m^3 and travel time of 1 s. It also means that pipe capacity is only applied to FTIP packages that are sent in each timestep. Thus, if a pipe has a pre-existing FTIP package from a previous timestep, that will not count toward the capacity calculation of the current timestep. This ensures that flow is accurate in steady state conditions but may not always be timestep accurate
2. An FTIP pipe may also back up based on a calculation performed on the pipe's outflow, given by Equation A4

$$V_{b,n} = \begin{cases} \sum_{m=0}^{n} V_{FTIP,m,i,j} - Q_{i,j}, & \sum_{m=0}^{n} V_{FTIP,m,i,j} > Q_{i,j} \\ 0, & \sum_{m=0}^{n} V_{FTIP,m,i,j} \leq Q_{i,j} \end{cases} \tag{A4}$$

where V_b is the volume of backed up water from FTIP package n that is about to be sent, $V_{FTIP,m,i,j}$ is the volume of a given FTIP package with a travel time of zero in pipe i,j, and m is either an FTIP package that has already been sent in that timestep or the package n, and $Q_{i,j}$ is the flow capacity of pipe i,j according to Equation A1 with units of volume per timestep.

Travel time in compartments is also tracked. A different method is used depending on whether the flow is arriving from an upstream compartment or for land runoff:

1. Flow arriving from upstream compartments is assumed to take the average full-bore travel time between all the sewer junctions in the given compartment that receive flow from upstream compartments and all the sewer junctions that send flow to downstream compartments
2. Flow from runoff is modeled using the time-area method (Butler & Davies, 2004). In pre-processing, when compartment-scale parameters are calculated, a time-area diagram is calculated for each compartment. This is normalized with total drainable surface area of a compartment to associate percentages of rainfall with travel times. The time-area method also assumes full-bore flow travel time calculations

Data Availability Statement

The software used in this study is available at (Dobson, Watson-Hill, et al., 2021). This can be considered for transparency purposes only, since the authors are currently creating a more user-accessible software implementation of CityWat-SemiDistributed for a wide range of integrated modeling purposes. The data used in this study is either provided at (ibid.) or has instructions for acquisition at (ibid.). The raw network data cannot be shared due to privacy reasons; however, we provide a version of the data with randomized locations and pipe/node parameters at (ibid.).

Acknowledgments
The research reported in this paper was taken as part of the CAMELLIA project (Community Water Management for a Liveable London), funded by the Natural Environment Research Council (NERC) under grant NE/S003495/1. The authors would like to thank Thames Water for providing the urban drainage model of Cranbrook. We are grateful to Yuting Chen for providing the high-resolution storm precipitation data used in this study.

References

Andrade, M. A., Choi, C. Y., Lansey, K., & Jung, D. (2016). Enhanced artificial neural networks estimating water quality constraints for the optimal water distribution systems design. *Journal of Water Resources Planning and Management*, *142*(9), 04016024. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000663

Babovic, F., & Mijic, A. (2019). The development of adaptation pathways for the long-term planning of urban drainage systems. *Journal of Flood Risk Management*, *12*(March), 1–12. https://doi.org/10.1111/jfr3.12538

Beven, K. (2012). *Rainfall-runoff modelling: The primer (2 nd ed.). Chichester: John Wiley and Sons*. https://doi.org/10.1002/9781119951001

Blondel, V. D., Guillaume, J. L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, *2008*(10), 1–12. https://doi.org/10.1088/1742-5468/2008/10/P10008

Bonald, T., de Lara, N., Lutz, Q., & Charpentier, B. (2020). Scikit-network: Graph analysis in Python. *Journal of Machine Learning Research*, *21*(185), 1–6.

Buluç, A., Meyerhenke, H., Safro, I., Sanders, P., & Schulz, C. (2016). Recent advances in graph partitioning. *Lecture notes in Computer Science (including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *9220 LNCS*, 117–158. https://doi.org/10.1007/978-3-319-49487-6_4

Burger, G., Bach, P. M., Urich, C., Leonhardt, G., Kleidorfer, M., & Rauch, W. (2016). Designing and implementing a multi-core capable integrated urban drainage modelling Toolkit: Lessons from CityDrain3. *Advances in Engineering Software*, *100*, 277–289. https://doi.org/10.1016/j.advengsoft.2016.08.004

Butler, D., & Davies, J. W. (2004). *Urban drainage* (2nd ed.). CRC Press.

DHI. (2014). *MOUSE pipe flow reference manual*. https://doi.org/10.1007/s13398-014-0173-7.2

Dobson, B., Jovanovic, T., Chen, Y., Paschalis, A., Butler, A., & Mijic, A. (2021). Integrated modelling to support analysis of COVID-19 impacts on London's water system and in-river water quality. *Frontiers in Water*, *3*(April), 26. https://doi.org/10.3389/frwa.2021.641462

Dobson, B., Watson-Hill, H., Muhandes, S., & Mijic, A. (2021). CityWat-SemiDistributed partition. https://doi.org/10.5281/zenodo.5095282

Falter, D., Vorogushyn, S., Lhomme, J., Apel, H., Gouldby, B., & Merz, B. (2013). Hydraulic model evaluation for large-scale flood risk assessments. *Hydrological Processes*, *27*(9), 1331–1340. https://doi.org/10.1002/hyp.9553

Gong, Y., Li, X., Zhai, D., Yin, D., Song, R., Li, J., & Yuan, D. (2018). Influence of rainfall, model parameters and routing methods on stormwater modelling. *Water Resources Management*, *32*(2), 735–750. https://doi.org/10.1007/s11269-017-1836-x

Hagberg, A. A., Schult, D. A., & Swart, P. J. (2008). Exploring network structure, dynamics, and function using NetworkX. In G. Varoquaux, T. Vaught, & J. Millman (Eds.), *Proceedings of the 7th Python in Science Conference* (pp. 11–15). Pasadena, CA, USA

Innovyze (2014). *InfoWorks ICM manual*. Portland

Karypis, G., & Kumar, V. (1997). *Metis: A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices*.

Liu, H., Zhao, M., Zhang, C., & Fu, G. (2018). Comparing topological partitioning methods for district metered areas in the water distribution network. *Water (Switzerland)*, *10*(4), 368. https://doi.org/10.3390/w10040368

Li, X., & Willems, P. (2020). A hybrid model for fast and probabilistic urban pluvial flood prediction. *Water Resources Research*, 1–26. https://doi.org/10.1029/2019wr025128

Muhandes, S., Dobson, B., & Mijic, A. (2021). A method for adjusting design storm peakedness to reduce bias in hydraulic simulations. *Proceedings of the Institution of Civil Engineers - Water Management* (pp. 1–13). Thomas Telford Ltd. https://doi.org/10.1680/jwama.20.00092

Ochoa-Rodriguez, S., Wang, L. P., Gires, A., Pina, R. D., Reinoso-Rondinel, R., Bruni, G., & Ten Veldhuis, M. C. (2015). Impact of spatial and temporal resolution of rainfall inputs on urban hydrodynamic modelling outputs: A multi-catchment investigation. *Journal of Hydrology*, *531*, 389–407. https://doi.org/10.1016/j.jhydrol.2015.05.035

Parés, F., Gasulla, D. G., Vilalta, A., Moreno, J., Ayguadé, E., Labarta, J., & Suzumura, T. (2018). Fluid communities: A competitive, scalable and diverse community detection algorithm. *Studies in Computational Intelligence*, *689*, 229–240. https://doi.org/10.1007/978-3-319-72150-7_19

Pedersen, A. N., Pedersen, J. W., Vigueras-Rodriguez, A., Brink-Kjær, A., Borup, M., & Mikkelsen, P. S., (2021). The Bellinge Data Set: Open Data and Models for Community-Wide Urban Drainage Systems Research. *Earth System Science Data Discussions*, *6*, 1–28. https://doi.org/10.5194/essd-13-4779-2021

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., & Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830

Razavi, S., Tolson, B. A., & Burn, D. H. (2012). Review of surrogate modeling in water resources. *Water Resources Research*, *48*(7), W07401. https://doi.org/10.1029/2011WR011527

Rossman, L. A. (2010). *Storm water management model user's manual*, version 5.0.

Salvadore, E., Bronders, J., & Batelaan, O. (2015). Hydrological modelling of urbanized catchments: A review and future directions. *Journal of Hydrology*, *529*(P1), 62–81. https://doi.org/10.1016/j.jhydrol.2015.06.028

Sela Perelman, L., Allen, M., Preis, A., Iqbal, M., & Whittle, A. J. (2015). Automated sub-zoning of water distribution systems. *Environmental Modelling & Software*, *65*, 1–14. https://doi.org/10.1016/j.envsoft.2014.11.025

Shi, J., & Malik, J. (2000). Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *22*(8), 888–905. https://doi.org/10.1109/34.868688

Thrysøe, C., Arnbjerg-Nielsen, K., & Borup, M. (2019). Identifying fit-for-purpose lumped surrogate models for large urban drainage systems using GLUE. *Journal of Hydrology*, *568*(July), 517–533. https://doi.org/10.1016/j.jhydrol.2018.11.005

Water, UK. (2019). *A framework for the production of drainage and wastewater management plans*.

Yan, D., Huang, L., & Jordan, M. I. (2009). Fast approximate spectral clustering. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 907–915). https://doi.org/10.1145/1557019.1557118