Network-Wide State Reconstruction for Urban Drainage Systems from Sparse Sensor Coverage

Imane El Ghabi[1], Revin Naufal Alief[2], Dilek Düştegör[2], Guy Henckens[3], Alexander Lazovik[2], and Victoria Degeler[1]

[1] Informatics Institute, University of Amsterdam, The Netherlands
i.elghabi@uva.nl
[2] Bernoulli Institute, University of Groningen, The Netherlands
[3] Aveco de Bondt, The Netherlands

Abstract. Comprehensive monitoring of urban drainage networks is essential for overflow mitigation and real-time operational decision-making, yet dense sensing remains economically impractical. We study network-wide hydraulic state reconstruction under sparse sensing by casting the drainage system as a heterogeneous graph and learning to infer unobserved node states from a small fraction of sensor observations. We propose a physics-informed heterogeneous attention graph neural network that injects hydraulic inductive bias into typed message passing (e.g., conveyance relationships and control-structure activation) while preserving computational efficiency. Using simulation data from a large municipal network in the Netherlands, we evaluate reconstruction performance at 5% sensor coverage (309 sensors over 6,441 nodes) across diverse rainfall scenarios. Results indicate that our approach achieves operationally relevant accuracy (MAE 0.432 m, ∼7% of the maximum operational depth (6 m)) with low inference latency (50–100 ms per network snapshot) suitable for real-time monitoring. We discuss practical implications and remaining challenges, including robustness to sensor noise and transfer across network topologies.

Keywords: urban drainage, sparse sensing, graph neural networks, state estimation, physics-informed learning

1 Introduction

Urban drainage operators require timely, network-wide situational awareness to prevent surcharge, flooding, and combined sewer overflows. However, real-world deployments typically instrument only a small subset of locations due to installation and maintenance costs. This motivates data-driven *spatial inference*: reconstructing hydraulic states at unobserved locations from sparse sensor observations.

Prior work has demonstrated the promise of graph neural networks for state estimation in wastewater contexts [1]. Building on this line, we focus on sparse-sensing reconstruction in large heterogeneous drainage networks and introduce a hydraulically structured inductive bias through typed message passing.

Sensor coverage. We define *sensor coverage* as the fraction of network nodes that provide measurements at inference time:

$$\text{coverage} = \frac{|\mathcal{V}_{\text{sensed}}|}{|\mathcal{V}|}, \tag{1}$$

where $\mathcal{V}$ is the set of nodes in the drainage graph and $\mathcal{V}_{\text{sensed}}$ is the subset of these nodes with sensors. In this paper we evaluate reconstruction at 5% coverage, meaning that 95% of the nodes must be inferred from sparse observations and network context.

Contributions. The main contributions of this work are summarized as follows:

- We formulate network-wide state reconstruction in urban drainage as learning on a *heterogeneous hydraulic graph* with typed nodes and edges.
- We introduce a physics-informed heterogeneous attention graph neural network that incorporates hydraulic structure into edge-type-specific message templates.
- We demonstrate reconstruction performance on a large municipal network under 5% sensing and discuss operational implications.

2 Network and Data

2.1 Hydraulic graph representation

We represent the drainage network as a directed heterogeneous graph $G = (\mathcal{V}, \mathcal{E})$. Nodes $\mathcal{V}$ correspond to hydraulic assets such as manholes, outfalls, and storage units. Edges $\mathcal{E}$ correspond to hydraulic connections such as conduits, pumps, and weirs (and other control structures when present). Typed nodes and edges capture distinct physical behavior and boundary conditions.

2.2 Simulation dataset

We use simulation-generated training and evaluation data from a Dutch municipal drainage network (Heusden, The Netherlands) with 6,441 nodes and 7,021 hydraulic connections. We generate 241 simulation scenarios spanning diverse rainfall patterns (e.g., design storms, historical events, and climate-stress tests) using InfoWorks ICM [8]. Each scenario provides time series of node-level hydraulic states (e.g., water depth), from which we treat each timestamp as an independent *snapshot* for training and evaluation.

2.3 Sensor placement and masking protocol

We evaluate reconstruction at 5% coverage by selecting a sensor set $\mathcal{V}_{\text{sensed}}$ and masking all other nodes at inference. Sensor placement is performed using K-means spatial clustering to maximize spatial coverage while ensuring uniform distribution across the network topology, yielding 309 instrumented nodes.

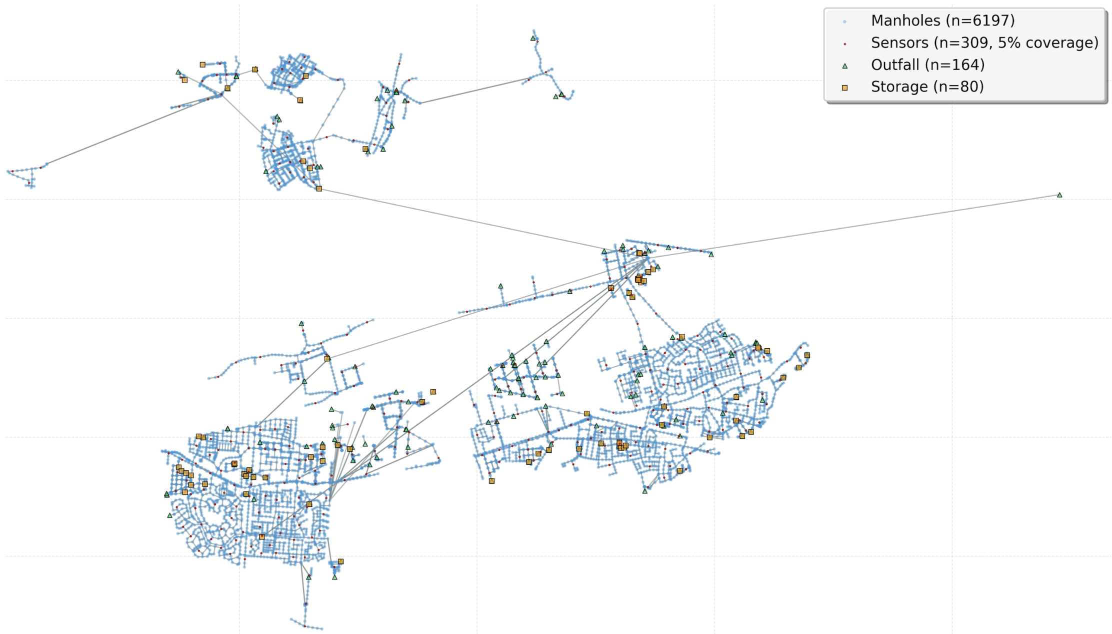

Fig. 1. Heusden drainage network as a heterogeneous graph. Nodes are colored by type (e.g., manholes/outfalls/storage). Sensed nodes at 5% coverage are highlighted.

3 Physics-Informed Heterogeneous Attention

3.1 Problem formulation

We consider state reconstruction on a fixed drainage graph $G = (\mathcal{V}, \mathcal{E})$. Each training/evaluation sample corresponds to a single simulated network *snapshot*. Let $x_v^{(s)}$ denote the target hydraulic state (e.g., water depth) at node v in snapshot s. At inference time, measurements $\tilde{x}_v^{(s)}$ are available only for $v \in \mathcal{V}_{\text{sensed}}$. The goal is to reconstruct $\hat{x}_v^{(s)}$ for all nodes $v \in \mathcal{V}$ given sparse observations and graph structure:

$$\{\hat{x}_v^{(s)}\}_{v\in\mathcal{V}} = f_\theta\Big(G,\ \{\tilde{x}_v^{(s)}\}_{v\in\mathcal{V}_{\text{sensed}}},\ \mathbf{a}_v,\ \mathbf{b}_e\Big), \tag{2}$$

where $\mathbf{a}_v$ are static node attributes and $\mathbf{b}_e$ are edge attributes.

3.2 Model overview

We propose a physics-inspired heterogeneous graph attention model designed to propagate information across drainage assets with different hydraulic roles. The network is represented as a heterogeneous graph with typed nodes (e.g., manholes, outfalls, storages) and typed edges (e.g., conduits, pumps, weirs). At inference time, only a sparse subset of nodes provides measurements (here: 5% coverage); the remaining node states are reconstructed by propagating information through the graph. Compared to our earlier work [1], which established the feasibility of GNN-based state estimation using homogeneous graph models, this method emphasizes heterogeneous, hydraulically motivated interactions aligned with drainage asset function under sparse sensing.

Typed message passing. As illustrated in the schematic in Figure 2, our model performs message passing where interactions are specialized by edge type. Hydraulically, distinct connection types govern flow dynamics differently: conduits convey gravity-driven flow, pumps reflect controlled transport, and weirs exhibit threshold-like activation. The network captures these behaviors through *edge-type-specific message templates* (Block A) that are conditioned on hydraulic attributes and combined across neighbors using attention-based weighting (Block B). Residual updates are used to stabilize multi-step propagation in large networks.

Physics-inspired inductive bias. Rather than enforcing full hydraulic equations during inference, the architecture introduces a hydraulically motivated inductive bias by structuring message templates according to asset function:

- **Conduits:** conveyance-oriented messages (Manning-inspired).
- **Weirs / control structures:** threshold-activation messages (crest-level-inspired).
- **Pumps:** controlled-transport messages (pump-curve-inspired).

Figure 2 provides a conceptual overview; detailed layer definitions and exact formulations are provided in an extended journal manuscript.

3.3 Input features and training

Node features include static asset attributes, sparse depth observations, and a sensor-availability flag. We train the model with a reconstruction loss over all nodes and all training snapshots:

$$\mathcal{L}(\theta) = \frac{1}{|\mathcal{S}_{\text{train}}|} \sum_{s \in \mathcal{S}_{\text{train}}} \left(\frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} \left| \hat{x}_v^{(s)} - x_v^{(s)} \right| \right), \tag{3}$$

with depth values normalized to zero mean and unit variance.

Curriculum masking. We employ a progressive masking schedule that starts with moderate masking (70% of nodes) and gradually increases sparsity toward 5% coverage (95% masking) during training. This curriculum learning approach improves stability under sparse supervision [5].

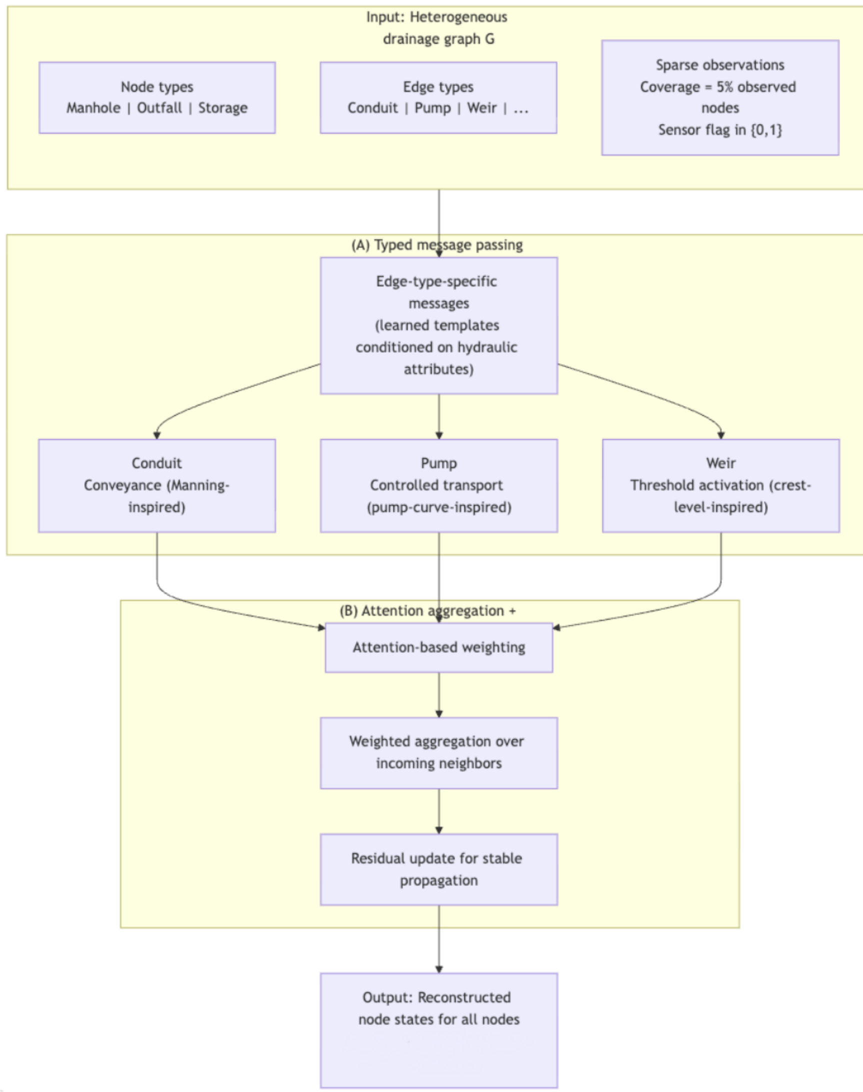

Fig. 2. Conceptual overview of the proposed model for sparse sensing. The drainage network is represented as a heterogeneous graph with typed nodes and edges. At inference time, only a small subset of nodes is observed (5% sensor coverage), indicated by a sensor flag. The model performs (A) edge-type-specific message passing using hydraulically motivated message templates (e.g., conveyance for conduits, controlled transport for pumps, and threshold activation for weirs), followed by (B) attention-based aggregation and residual updates to reconstruct node states across the full network. Detailed layer definitions and exact formulations are provided in an extended journal manuscript.

4 Experimental Setup

We split simulations into training/validation/test sets using a 60/20/20 ratio, ensuring that scenarios from the same rainfall event remain within the same split to prevent data leakage. Performance is reported using mean absolute error (MAE) averaged over test snapshots:

$$\mathrm{MAE} = \frac{1}{|\mathcal{S}_{\text{test}}|} \sum_{s \in \mathcal{S}_{\text{test}}} \left(\frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} \left| \hat{x}_v^{(s)} - x_v^{(s)} \right| \right). \tag{4}$$

Baselines. We compare our approach against (i) a homogeneous GAT variant that ignores edge/node types and treats all connections uniformly [3], and (ii) a spatial interpolation baseline (inverse distance weighting) applied on the graph topology.

5 Results and Discussion

At 5% sensor coverage, the proposed model achieves MAE 0.432 m across the test set, representing $\sim$7% of the maximum operational depth (6 m). Inference latency is low (50–100 ms per network snapshot) relative to full hydraulic simulation, supporting near-real-time monitoring. Comparisons with the homogeneous GAT baseline reveal that the proposed heterogeneous architecture significantly reduces reconstruction error, validating the inclusion of hydraulic inductive biases.

Transferability. Our evaluation in this work is conducted on a single network graph (Heusden). To ensure generalizability, future work will evaluate the model on other network topologies. Specifically, the architecture will be trained and evaluated separately on entirely distinct datasets (for example, contrasting closed municipal networks in the Netherlands with the open Shunqing dataset in China) rather than being adapted from one to the other, which will rigorously test the model's robustness across different asset mixes and connectivity patterns.

Limitations. Current experiments assume simulation-consistent measurements and do not fully capture real sensor noise and outages. Real deployments can exhibit 1–5% measurement error and intermittent missingness; robustness to these factors will be addressed in extended evaluation.

6 Conclusion

We presented a physics-informed heterogeneous attention graph neural network for network-wide state reconstruction in urban drainage systems under sparse sensing. Experiments on a large municipal network demonstrate promising performance at 5% coverage and computational efficiency suitable for operational settings. Future work will focus on robustness to measurement noise and transferability across network topologies.

Acknowledgements This research is funded by the Dutch Research Council (NWO) Merian Fund – Cooperation China (CAS) program through the project DDT-Clean: Intelligent Wastewater Treatment – Distributed Digital Twin for Clean Water (482.22.007).

References

1. El Ghabi, I., Naufal, R., Düştegör, D., Henckens, G., Lazovik, A., Degeler, V.: GNN-Based State Estimation for Urban Wastewater Digital Twins. In: Proceedings of the 41st ACM/SIGAPP Symposium on Applied Computing (SAC), 2026. doi: 10.1145/3748522.3779823
2. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: International Conference on Learning Representations (ICLR) (2017)
3. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. In: International Conference on Learning Representations (ICLR) (2018)
4. Hamilton, W.L.: Graph Representation Learning. Synthesis Lectures on Artificial Intelligence and Machine Learning, 14(3), 1–159 (2020)
5. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: Proceedings of the 26th International Conference on Machine Learning (ICML), pp. 41–48 (2009)
6. Sit, M., Demiray, B.Z., Xiang, Z., Ewing, G.J., Sermet, Y., Demir, I.: A comprehensive review of deep learning applications in hydrology and water resources. Water Science and Technology, 82(12), 2635–2670 (2020)
7. Truong, H., Tello, A., Lazovik, A., Degeler, V.: Graph Neural Networks for Pressure Estimation in Water Distribution Systems. Water Resources Research, 60(4), e2023WR035123 (2024)
8. InfoWorks ICM (Innovyze / Autodesk): Integrated catchment modelling software. https://www.autodesk.com/products/infoworks-icm/overview (accessed: 2024-12-15)