GHOST IN TOPOLOGICAL NEURAL FLUX PREDICTION: BOUNDARY CONDITIONS MATTER

Anonymous authors

000

001

002003004

010 011

012

013

014

015

016

017

018

019

021

024

025

026

027

028

029

031

033

035

037

040

041

042

043

044

045

047

048

051

052

Paper under double-blind review

ABSTRACT

Topological flux prediction (TFP) aiming to model spatiotemporal fluid transport over networked systems, has inspired and lent itself to various predictive methods. Whereas Graph Neural Networks (GNNs) demonstrate successes in related prediction tasks, recent studies suggest that they can underperform even simple baselines in TFP, concluding that GNNs may be ill-suited for such problems. In this paper, we re-examine this claim by dissecting the learning behavior of GNNs on fluid networks, decoupling the roles of boundary nodes, which regulate total influx, from interior nodes. We find that the dominant prediction errors arise at boundary nodes, which do not necessarily imply a fundamental limitation in the expressive power of GNNs. We interpret this phenomenon from a dynamical-systems perspective, arguing that GNNs incur substantial boundary losses mainly due to the lack of explicit modeling of boundary conditions. To compensate this information deficit, we propose a novel ghost-TFP framework, which learns ghost node proxies with an implicit solver to capture boundary-aware representations. Experimental results on two real datasets show that our method ghost-TFP improves standard GNNs by reducing the average MSE by 8.35% and 5.0%, and the boundary node MSE by 11.2% and 7.1%, respectively. For efficiency, we further devise an explicit solver that learns inverse operators which, depending on the underlying GNN backbone, can accelerate inference by $2\times$ on both datasets. Codes are available at https://anonymous.4open.science/r/Ghost-Node-GNN-FB1E.

1 Introduction

Flux prediction enjoys broad use in fluid systems (Kratzert et al., 2021; Jin et al., 2023), *e.g.*, flood forecasting (Jiang et al., 2025; Bentivoglio et al., 2025b), hydrochemical modeling (Mangold & Tsang, 1991), estuarine circulation (Geyer & MacCready, 2014), among others. Since such systems unfold in both space and time, Graph Neural Network (GNN) (Zhou et al., 2020; Wu et al., 2020) emerges as a seemingly plausible modeling choice, given its demonstrated success in related tasks such as traffic forecasting (Jin et al., 2023) and energy transmission (Varbella et al., 2024).

Yet, recent studies reveal that GNNs often ignore the underlying fluid dynamics and learn absurd patterns from data, *e.g.*, predicting fluxes moving from downstream to upstream, which violate gravity and conservation laws (Kirschstein & Sun, 2024). As a result, some conclude that incorporating fluid system topology offers little benefit, as GNNs underperform even simple baselines like multilayer perceptrons (MLPs), which do not model graph structure at all (Kirschstein & Sun, 2024).

In this paper, we argue that such conclusions are premature. We revisit the prediction loss patterns from GNN-based flux models and find that the dominant errors lie at the *boundary nodes*. As Figure 1 illustrates, whereas boundary nodes account for more than half of the total nodes in a fluid network, existing models consistently yield much higher prediction errors on them compared to interior nodes. In fact, if we restrict evaluation to interior nodes alone, GNN predictors significantly outperform baseline MLPs. This discrepancy suggests that treating all nodes equally without distinguishing boundary from interior may underlie the perceived failure of GNNs in fluid systems.

To scrutinize this observation further, we analyze the GNN learning behavior from a dynamical-system perspective (Poli et al., 2019), where message-passing simulates the update of node states

Avg. MSE per node

Model	Boundary	Interior	Diff. (%)
ResGCN	0.1408	0.0985	30.04%
ResGAT	0.1401	0.0836	40.33%
GCNII	0.1511	0.0891	41.03%
GNN error (avg.)	0.12	17	-
MLP	0.1210	0.1030	14.89%
MLP error (avg.)	0.11	35	-

Figure 1: A fluid network from the LamaH-CE2 dataset, where boundary nodes account for more than half of the graph topology. Although standard GNNs outperform a simple baseline (e.g., MLP) on interior nodes, a misleading conclusion may arise due to their larger loss on boundary nodes, resulting in a higher overall average MSE. We hypothesize that this is because GNNs lacking explicit modeling of boundary condition, which is essential in classic PDE-based flux prediction models for networked fluid systems.

under local interaction rules (*e.g.*, a differential equation), while each such layer proceeds a discrete time step (*e.g.*, an Euler step). Boundary conditions in such systems constrain their solution space (LeVeque, 2007). In GNN-based flux models, boundary nodes are discrete counterparts of these conditions, as they regulate influx into the entire network. We hypothesize that such regulation power of boundary conditions acts as strong inductive bias, and incorporating it into the learning process is critical to unleashing the full potential of GNNs in flux prediction tasks.

To instantiate and validate our hypothesis, we propose a novel computing paradigm, termed <code>ghost-TFP</code>, where we borrow the ghost node technique (Tseng & Ferziger, 2003) from finite-difference methods (FDM) to estimate boundary conditions. Specifically, we approximate ghost nodes by extrapolating from boundary nodes and their immediate (downstream) neighbors. This imposes recursive coupling on message-passing, where each boundary node depends on ghost nodes, whose embeddings are in turn defined by interior and boundary nodes themselves. Unknown node representations thus appear on both sides of the update equations, making standard layer-wise GNN training inapplicable. To solve this, we adopt implicit GNN (Gu et al., 2020), which recasts message-passing as a fixed-point problem and seeks an equilibrium that aligns node representations with the structural relationships determined by boundary conditions.

Note, imposing different boundary conditions on fluid system may derive disparate structural couplings among ghost, boundary, and interior nodes. Each such coupling defines a unique augmentation of graph adjacency, which entails extensive craftsmanship to adapt the implicit GNN solver to every possible augmentation. To counter this, we further unify our framework by treating the inverse of the augmented adjacency as a learnable operator, lending a closed-form approximation to the implicit solution. **Specific contributions** in this paper are summarized as follows.

- (1) We decouple error sources in GNN-based flux models and find the dominant loss stems from missing explicit boundary-node modeling.
- (2) We propose ghost-TFP to learn boundary conditions via ghost-node proxies; it outperforms GNN base-lines by 6.68% on average and mitigates the boundary-interior loss gap by 11.03% on two real datasets.
- (3) We devise a unified operator-learning view within ghost-TFP that avoids hand-crafted adjacency in implicit solvers, yielding up to 2× training speedup (backbone-dependent) without sacrificing accuracy.

2 Preliminaries

Notation and Problem Statement. Let G=(V,E) denote a directed fluid network, where nodes V represent local observation points and edges E indicate the direction of flow. We define the adjacency $\mathbf{A} \in \{0,1\}^{|V| \times |V|}$, such that $\mathbf{A}_{i,j} = 1$ if there exists a directed edge from node v_i to v_j , and $\mathbf{A}_{i,j} = 0$ otherwise. Note that $\mathbf{A} \neq \mathbf{A}^{\top}$. A node v_j is said to be a neighbor of v_i , denoted $v_j \in \mathcal{N}(i)$, if $\mathbf{A}_{j,i} = 1$.

At time t, each node $v_i \in V$ is associated with a feature matrix $\mathbf{h}_i \in \mathbb{R}^{W \times d}$, which stores d physical measurements (e.g., velocity, pressure, slope) over a historical window of W time steps. Stacking features across all nodes yields a tensor $\mathbf{H} = [\mathbf{h}_1, \dots \mathbf{h}_{|V|}]^{\top} \in \mathbb{R}^{|V| \times W \times d}$. The goal of flux prediction is to learn a predictive model f that forecasts a target physical quantity (e.g., flux volume) at a future step f and f the prediction horizon. Let $\mathbf{y} \in \mathbb{R}^{|V|}$ denote the ground-truth values of this target quantity. Our learning objective is to minimize the empirical loss f (\mathbf{y} , f (\mathbf{H} , \mathbf{A})).

We define the *boundary nodes* $V_{\rm BN}$, as those with zero in-degree, *i.e.*, $\deg^-(v_b) = 0$, $\forall v_b \in V_{\rm BN}$. Intuitively, these nodes lie at the most upstream points of G and regulate the influx into the whole system. The remaining nodes are referred to as *interior nodes*, defined as $V_{\rm IN} = V \setminus V_{\rm BN}$.

3 PROBLEM ANALYSIS

A paradoxical observation is that if the predictive model f is instantiated as GNN, it often underperforms simple baselines (e.g., MLP) and shows little difference between predictions computed from $f(\mathbf{H}, \mathbf{A})$ and $f(\mathbf{H}, \mathbf{A}^{\top})$. Here, as \mathbf{A} encodes the ground-truth forward fluid flow, \mathbf{A}^{\top} represents a physically-implausible, reversed flow. Kirschstein & Sun (2024) concludes that GNNs may not work in flux prediction tasks, as they fail to distinguish directional flow consistency, which is fundamental in physical systems.

3.1 GNNs as Neural Differential Equations.

We argue that such a conclusion is premature and advocate an alternative interpretation through the lens of numerical solvers (Liu et al., 2025). Write a standard message-passing update as

$$\mathbf{h}_{i}^{(l+1)} = \mathbf{h}_{i}^{(l)} + \sum_{v_{i} \in \mathcal{N}(i)} \psi^{(l)}(\mathbf{h}_{i}^{(l)}, \mathbf{h}_{j}^{(l)}), \ \mathbf{h}_{i}^{(0)} = \mathbf{h}_{i},$$
(1)

where $\mathbf{h}_i^{(l)}$ denotes the representation of node v_i at layer l, and $\psi^{(l)}$ is the message function aggregating information from its upstream neighbors $v_j \in \mathcal{N}(i)$. We view Eq. (1) as an explicit integration scheme, that mimics a partial differential equations (PDEs) solver (LeVeque, 2007), such as

$$u^{t+1}(x_i) = u^t(x_i) + \Delta t \cdot F(u^t(x_i), u^t(x_{i+1})), \tag{2}$$

s.t.
$$u^t(x_{i+1}) \approx B(u^t(x_i)), \forall v_i \in V_{BN}$$
 (3)

where $u^t(x_i)$ is the state of a physical quantity (e.g., flux volume) at location x_i and time t, and \mathcal{F} represents the spatial derivative or transport dynamics. Consider, for example, a discretized advection equation (Chock, 1991) that implements $F(u^t(x_i), u^t(x_{i+1})) = \frac{\partial}{\partial x}(u^t(x_{i+1}) - u^t(x_i))$, which models the propagation of the quantity through its upstream neighbor x_{i+1} . We can observe an algebraic similarity between Eq. (1) and Eq. (2), where the GNN layer index t parallels the time step t, and the message function ψ acts as the discrete derivative F across the graph topology.

A natural question arises: if PDE solvers operate robustly and in a topology-aware manner for fluid systems, and message-passing mimics their computational structure, then *why do GNNs empirically fail on the same tasks?*

Boundary Information Deficit. We hypothesize that a critical missing component in message-passing is the lack of boundary information. PDE solvers explicitly enforce such conditions, as seen in Eq. (3), which defines the derivative term F at the boundary location x_b . Intuitively, when x_b lies at the spatial boundary, its upstream neighbor x_{b+1} is undefined, invalidating the computation of the flux term F in Eq. (2). The boundary condition in Eq. (3) complements this by prescribing how F should be computed solely from the local state $u^t(x_b)$. To wit, a Robin-type (Busse et al., 2017) boundary condition specifies Eq. (3) as

$$u^{t}(x_{b+1}) = \omega_1 \cdot u^{t}(x_b) + \omega_2 \cdot \partial u^{t}(x_b) / \partial x, \tag{4}$$

with $\omega_1, \omega_2 \in \mathbb{R}$, which closes the dynamical system by postulating an interpolation between x_b and its spatial derivative to approximate the undefined, out-of-boundary x_{b+1} .

In contrast, GNNs lack mechanism to compensate for this information deficit at graph boundaries. For a boundary node v_b which, by definition, has no incoming edge and thus no upstream neighbor, namely $\mathcal{N}(b) = \emptyset$. As a result, Eq. (1) collapses to $\mathbf{h}_b^{(l+1)} = \mathbf{h}_b^{(l)} + \psi^{(l)}(\mathbf{h}_b^{(l)}, \mathbf{0})$, meaning that the update of v_b depends solely on its own features and receives no information from the graph topology. This isolation over successive layers leads to degraded boundary node embeddings and, eventually, to substantial prediction errors.

3.2 EMPIRICAL VALIDATION.

To validate our hypothesis, we analyze and compare the prediction losses of ResGAT (Residual Graph Attention Networks) and MLP by separating the errors incurred at boundary versus interior nodes. The results are summarized in Table 1.

Table 1: Node-wise MSE Comparison

We make two observations from these results. *First*, the overall mean squared error (MSE) misleadingly suggests that the GNN underperforms an MLP (.1166 > .1135), when in fact the GNN performs substantially better in regions where it can leverage

Node Type	MSE For. (A)	MSE Rev. (\mathbf{A}^{\top})
Boundary (V_{BN}) Interior (V_{IN}) All nodes (V)	.1401 .0836 .1166	.1350 .0939 .1179
MLP	.1	135

graph topology. The failure of GNN is mainly attributed to the boundary nodes, incurring a .1401

163 164

166 167

168 169 170

171

172

173

174

175

176

177

178

179

181

182

183

185

186

187

188

189

190 191

192

193

194

195

196

197

199

200

201

202

203

204

205 206

207

208 209

210

211 212

213

214 215

Figure 2: Overview of the proposed ghost-TFP framework. Left: Ghost-node construction (Section 4.1). For each boundary node v_b , we introduce a corresponding ghost node v_q . Its embedding h_q is learned from the boundary embedding h_b and the aggregated embeddings of its downstream interior neighbors h_{nbr} . Middle: Implicit solver (Section 4.2). We perform implicit message passing on the augmented graph with ghost nodes using a fixed-point solver, converging to \mathbf{H}'^{\star} . Shared parameters Θ are applied consistently across updates. Right: Explicit inverse-operator learning (Section 4.3). We learn an explicit inverse operator on the densified adjacency $\mathbf{A}_{\mathbf{\Theta}}^{\prime}$, enabling layer-wise updates through a learnable inverse mapping.

MSE. This large boundary loss obscures the otherwise strong performance of GNN-based flux prediction on interior nodes, where MSE drops to .0836. Second, the seemingly similar overall losses using the ground-truth forward flow A (.1166) and the reverse flow A^{\top} (.1179) are deceptive. The forward model is disproportionately penalized by boundary node errors, which suppress its average performance and mask its superiority over the reversed model. When focusing on interior nodes only, the forward model achieves an MSE of .0836, outperforming the reversed model at .0939. This aligns with physical intuition and demonstrates the positive impact of graph topology on learning meaningful representations in regions where directional flow information is available.

Note, these MSE results are normalized, where .001 change means 1\% flux volume change. For a mid-size river network (Discharge: 100m³/s), a .01 MSE error in discharge prediction equals a daily volume discrepancy of 86,400 m³, which equals to \sim 35 Olympic pools. This may cause critical failure with cascading consequences, e.g., threaten the survival of aquatic species (Poff et al., 1997), lead to dangerous underestimations of pollution risk (Whitehead et al., 2009), and be amplified into financial losses for hydropower and navigation (Lehner et al., 2005; Jonkeren et al., 2007).

THE GHOST-TFP APPROACH

This section presents our ghost-TFP computing paradigm. Traditional GNNs lack the mechanism to model upstream boundary influence, which leads to degraded performance near the system boundaries. To address this, we introduce ghost nodes as boundary-aware proxies that restore physical consistency in both implicit and explicit settings. In Section 4.1, we describe how to construct and learn ghost node embeddings based on surrounding interior and boundary information. Section 4.2 formulates an implicit solver that integrates ghost nodes into a fixed-point framework inspired by numerical PDE methods. Section 4.3 further provides a computationally efficient, explicit counterpart by learning a dense inverse operator, allowing layer-wise updates that preserve boundary-awareness while avoiding the cost of fixed-point iteration.

4.1 LEARN GHOST NODE PROXIES

To remedy the boundary information deficit in GNNs, we propose to learn the boundary proxy from data through the ghost node method (Tseng & Ferziger, 2003). Let v_q denote a virtual ghost node corresponding to a boundary node v_b , and v_{nbr} denote the interior (downstream) neighbor of v_b , such that $v_{\text{nbr}} \in \{v_j \mid v_b \in \mathcal{N}(j)\}$. We draw insights from PDE solvers to learn the embedding of v_q from the structural coupling among v_g , v_b , and v_{nbr} . Specifically, we can discretize a Robin-type boundary condition in Eq. (3) and Eq. (4) to derive $\mathbf{h}_g^{(l)} = \omega_1 \cdot \mathbf{h}_b^{(l)} + \omega_2 \cdot (\mathbf{h}_b^{(l)} - \mathbf{h}_{nbr}^{(l)})/\Delta x$, where $\mathbf{h}_{\mathrm{nbr}}^{(l)}$ is the embedding of v_{nbr} at the l-th layer. We parameterize ω_1 and ω_2 via an MLP, defined as $\mathbf{h}_g^{(l)} = \mathrm{MLP}(\mathrm{Concat}(\mathbf{h}_b^{(l)}, \mathbf{h}_{\mathrm{nbr}}^{(l)}); \theta_{\mathrm{gh}}).$

$$\mathbf{h}_g^{(l)} = \text{MLP}(\text{Concat}(\mathbf{h}_b^{(l)}, \mathbf{h}_{\text{nbr}}^{(l)}); \theta_{\text{gh}}). \tag{5}$$

The total number of such ghost nodes equates to boundary nodes. Defining the set of ghost nodes $V_{\rm GH} = \{v_q\}$, we have $|V_{\rm GH}| = |V_{\rm BN}|$. To proceed message-passing, we define a graph augmentation operator \mathcal{A} , which takes the original graph as inputs and augments it with ghost nodes as follows.

$$(\mathbf{A}', \mathbf{H}') = \mathcal{A}(\mathbf{A}, \mathbf{H}),\tag{6}$$

where $\mathbf{H}' = [\mathbf{H}, \{\mathbf{h}_g\}_{|V_{\mathrm{BN}}|}]^{\top} \in \mathbb{R}^{(|V|+|V_{\mathrm{GH}}|) \times W \times d}$ denotes the augmented node feature matrix, and $\mathbf{A}' \in \{0,1\}^{(|V|+|V_{\mathrm{GH}}|) \times (|V|+|V_{\mathrm{GH}}|)}$ is the augmented adjacency, such that $\mathbf{A}'_{i,j} = 1$ if v_i is an (upstream) neighbor of v_j , namely $v_i \in \mathcal{N}(j)$, and $\mathbf{A'}_{i,j} = 0$ otherwise.

217

218

219

220 221

222

224

225 226

227

228

229

230 231

232 233

234

235 236

237 238

239

240

241

242 243

244 245

246

247

249

250

251

252 253

254

255 256

257

258

259

260

261

262

263 264

265

266 267

268

269

The augmentation in Eq. (6) connects each ghost node to a downstream boundary neighbor in a deterministic way, which challenges layer-by-layer message passing. Specifically, the update of ghost nodes in numerical methods (Tseng & Ferziger, 2003) takes the following form

$$(1 - \frac{\delta_1^2 \Delta t}{\delta_2}) u^{t+1}(x_b) + (\frac{\delta_1 \Delta t}{\delta_2}) u^{t+1}(x_g) = u^t(x_b), \tag{7}$$

 $(1-\frac{\delta_1^2\Delta t}{\delta_2})u^{t+1}(x_b)+(\frac{\delta_1\Delta t}{\delta_2})u^{t+1}(x_g)=u^t(x_b), \tag{7}$ where $\delta_1,\delta_2\in\mathbb{R}$ are two physical coefficients. The derivation from Eq. (2) and Eq. (3) to Eq. (7) are deferred to Appendix 8.5 due to the space limit. Drawing analogy to Eq. (7), learning the ghost node embeddings is constrained by $\alpha_1 \cdot \mathbf{h}_b^{(l+1)} + \alpha_2 \cdot \mathbf{h}_g^{(l+1)} = \mathbf{h}_b^{(l)}, \ \exists \alpha_1, \alpha_2 \in \mathbb{R}$, resulting in the message-passing on boundary node as:

$$\mathbf{h}_{b}^{(l+1)} = \mathbf{h}_{b}^{(l)} + \sum_{v_{g} \in \mathcal{N}(b)} \psi^{(l)}(\mathbf{h}_{b}^{(l)}, \mathbf{h}_{g}^{(l)}), \quad \text{s.t. } \alpha_{1} \cdot \mathbf{h}_{b}^{(l+1)} + \alpha_{2} \cdot \mathbf{h}_{g}^{(l+1)} - \mathbf{h}_{b}^{(l)} \approx 0.$$
 (8)

Main steps for ghost-node learning are presented in Algorithm 8.7.1 in Appendix 8.7. Note that computing $\mathbf{h}_q^{(l+1)}$ again requires information from $\mathbf{h}_h^{(l+1)}$ and its downstream neighbor, as indicated by Eq. (5). This forms a coupled system (LeVeque, 2007), where unknown variables appear on both sides of the update equation, necessitating implicit solver via fixed-point iteration.

4.2 IMPLICIT GHOST-BOUNDARY MESSAGE-PASSING

To learn Eq. (8), we propose an implicit GNN solver, aiming to find an equilibrium that satisfies

$$\min_{f,\theta_{\text{sh}},\Theta} \ell(\mathbf{y}, f(\mathbf{H}'^{(L)})), \quad l = [0, \dots, L-1], \tag{9}$$

s.t.
$$\Pi_{\mathbf{A}}(\mathcal{I}(\mathbf{A}'\mathbf{H}'^{(l+1)}\Theta)) = \mathbf{H}^{(l)}, \ \|\Theta\|_{\infty} \le 1/\lambda_{pf}(\mathbf{A}'),$$

where $\mathcal{I}(\cdot)$ represents the implicit solver. $\Pi_{\mathbf{A}}(\cdot)$ projects node embeddings from the augmented graph (including ghost nodes learned by θ_{gh}) back to the original node indices in G. Unlike traditional layer-wise message-passing, where each layer l is associated with its own learnable weight matrix, our design shares one single parameter matrix Θ across all implicit update steps (i.e., between l and l + 1). We adopt the fixed-point update scheme (Gu et al., 2020) to solve Eq. (9).

The gradient of prediction loss ℓ with respect to the parameter $W \in \theta_{\mathrm{gh}} \cup \Theta$ can be derived as $\nabla_W \ell = \langle \partial (\mathbf{A}' \mathbf{H}'^{(l+1)} \Theta) / \partial W, \nabla_Z \ell \rangle$, where $\nabla_Z \ell$ is the derivative through a fixed point X that satisfies the implicit relation $Z = \sigma(\mathbf{A}'X\Theta)$, with the non-linear activation σ embedded in $\mathcal{I}(\cdot)$ for improved representational expressivity (Chen et al., 2023). This fixed point induces an implicit gradient $\nabla_Z \ell = D \odot (\mathbf{A}' \ \nabla_Z \ell \Theta^\top + \nabla_X \ell)$, where $D = \sigma'(\mathbf{A}' X \Theta)$ is the Jacobian of the activated output, computed element-wise as $\sigma'(z) = d\sigma(z)/dz$, and \odot denotes an element-wise multiplication. The equilibrium for $\nabla_z \ell$ enjoys a unique solution under well-posedness conditions by iterating it to convergence. In practice, we adapt the Picard iteration (Berinde, 2007) to jointly compute $\nabla_{\mathbb{Z}}\ell$ and X at each iteration. Once we have $\nabla_Z \ell$, the gradient $\nabla_W \ell$ immediately follows by using chain rule via auto-differentiation (Ren et al., 2023).

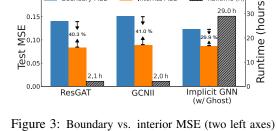
To ensure the unique solution of Θ , we follow (Gu et al., 2020) to impose a spectral constraint via Perron-Frobenius norm. Intuitively, if the spectral strength $\mathbf{A}'^{\top}\Theta$ is too large, the fixed-point iteration may diverge. To analyze and control its magnitude, we vectorize the implicit mapping as $\text{vec}(\mathbf{A}'\mathbf{H}'^{(l+1)}\Theta) = (\Theta^{\top} \otimes \mathbf{A}') \text{ vec}(\mathbf{H}'^{(l+1)}), \text{ so that we can single out the impact of } \mathbf{A}'^{\top}\Theta \text{ during}$ optimization. The Kronecker product \otimes multiplies each element of A' onto the matrix Θ (Schacke, 2004), which expands the bilinear operator into a matrix acting on the vectorized node embeddings. The Perron-Frobenius eigenvalue $\lambda_{pf}(\mathbf{A}')$ computes the largest eigenvalue of the augmented adjacency A', which measures the maximal amplification factor of the augmented graph topology (Berman & Plemmons, 1994). $\|\Theta\|_{\infty}$ takes the maximum absolute row sum, which provides a convex upper bound on its spectral norm (Zheng & Wang, 2008). To stabilize implicit updates and expedite convergence, we have $\lambda_{pf}({\mathbf{A}'}^{\top}\Theta) \leq \lambda_{pf}({\mathbf{A}'})\|\Theta\|_{\infty} \leq 1$ to bound the spectral strength $\mathbf{A}^{\prime} \Theta$ which gives the constraint term in Eq. (9). Algorithm 8.7.2 in Appendix 8.7 summarizes the main steps for the implicit fixed-point solver.

4.3 EXPLICIT ADJACENCY-INVERSE SOLVER

While the ghost nodes compensate for the boundary information deficit, and Eq. (9) provides an implicit yet effective solution for them, this solution suffer from two key efficiency limitations. First,

as shown in Figure 4a, although the ghost node-enhanced implicit GNN reduces the prediction loss on boundary nodes by 7.2 %, it suffers from high computational cost, running approximately $13 \times$ slower than standard layer-wise message-passing GNNs. Second, the structure of the augmented adjacency \mathbf{A}' can vary w.r.t. the modeling choice of boundary condition. In practice, different numerical schemes, e.g., Robin-type as we used in Eq. (4), can be used to impose boundary constraint on the same physical system. As such, our implicit solver that assumes a fixed \mathbf{A}' loses flexibility in adapting to such variations to learn ghost node proxies.

Observing Eqs. (5), (6), and (8), we note that although \mathbf{A}' connects each ghost node v_g only to its downstream boundary node v_b , the computation of \mathbf{h}_g in fact depends on both \mathbf{h}_b and the interior neighbor embeddings \mathbf{h}_{nbr} . This reveals an implicit computational graph that spans \mathbf{h}_g , \mathbf{h}_b , and \mathbf{h}_{nbr} , lending to consider whether we can construct a weighted adjacency matrix $\mathbf{A}'_{\mathbf{\Theta}} \in \mathbb{R}^{(|V|+|V_{GH}|)\times(|V|+|V_{GH}|)}$ that captures these interactions in a layer-wise message-passing regime.



and wall-clock runtime (right axis) for two backbone

r-wise GNNs and our proposed implicit solver with ghost nodes. Observe that while the implicit solver reduces prediction error at boundary nodes and mitigates the boundary-interior loss gap, it incurs a runtime overhead of over 14× compared to standard GNNs

To this end, we decompose the augmented adjacency \mathbf{A}' into two subgraphs. The *first* is a of over $14 \times$ compared to standard GNNs.

standard GNN over the augmented node set $V \cup V_{\rm GH}$, where each node (whether interior, boundary, or ghost) participates in message passing. We let $w_{i,j}$ denote the message-passing weight from node v_j to v_i , with $v_i, v_j \in V \cup V_{\rm GH}$. The **second** is a bipartite subgraph linking ghost nodes $V_{\rm GH}$ to original nodes $V_{\rm CH}$, which reflects how ghost nodes are constructed from the boundary condition. Let $p_{i,j}$ denote the interpolation weight from an original node $v_i \in V$ to a ghost node $v_j \in V_{\rm GH}$, of which the edges are restricted to connect nodes across the bipartition. This allows us to encode various boundary condition types in a unified message-passing framework. For example, write \mathbf{h}_g and \mathbf{h}_b the ghost and boundary node embeddings, respectively, and let $\mathbf{h}_{\mathrm{nbr}_1}$, $\mathbf{h}_{\mathrm{nbr}_2}$ denote the embeddings of the immediate and second-order interior (downstream) neighbors of boundary, respectively. Under a first-order condition, the ghost node proxy is defined as $\mathbf{h}_g = p_{b,g}\mathbf{h}_b + p_{n_1,g}\mathbf{h}_{\mathrm{nbr}_1}$. For second-order, it becomes $\mathbf{h}_g = p_{b,g}\mathbf{h}_b + p_{n_1,g}\mathbf{h}_{\mathrm{nbr}_1} + p_{n_2,g}\mathbf{h}_{\mathrm{nbr}_2}$. Each such augmentation changes the structure of the rows in \mathbf{A}' corresponding to v_g , introducing new learnable parameters in \mathbf{A}'_{Θ} . Upon these intuitions, we define (i,j)-th entry of \mathbf{A}'_{Θ} as follows.

$$\mathbf{A'_{\Theta}}[i,j] = \begin{cases} 1 & \text{if } i = j \\ p_{i,j} & \text{if } v_i \in V \text{ and } v_j \in V_{\text{GH}} \\ w_{i,j} & \text{if } v_i \in V \text{ and } v_j \in \mathcal{N}(i) \\ 0 & \text{otherwise} \end{cases}.$$

This \mathbf{A}'_{Θ} reduces the implicit solver in Eq. (9) to an explicit solution linear system $\Pi_{\mathbf{A}} \left(\mathbf{A}'_{\Theta} \mathbf{H}'^{(l+1)} \right) = \mathbf{H}^{(l)}$, which enjoys a closed-form solution $\mathbf{H}'^{(l+1)} = \mathbf{A}'_{\Theta}^{\dagger} \mathbf{H}'$. Here, \dagger denotes the Moore-Penrose inverse (Prasad & Bapat, 1992), and the ghost node-padded tensor \mathbf{H}' is defined in Eq. (6). We present a concrete examples in Appendix 8.6. One key property of the inverse is producing dense matrices (Chamberlain et al., 2021), which reflects the global coupling of the system and allows to bypass the limitation of using a sparse, hand-crafted \mathbf{A}' . To operationalize this idea, we learn $\mathbf{A}'_{\Theta}^{\dagger}$ by approximating the inverse operator through a trainable GNN as

$$\min_{f,\theta_{\text{gh}},\Psi} \ell(\mathbf{y}, f(\mathbf{A}_{\Theta}^{\prime}^{\dagger}\mathbf{H}^{\prime(L)})), \quad l = [0, \dots, L-1],$$
s.t.
$$\mathbf{H}^{\prime(l+1)} = \mathbf{A}_{\Theta}^{\prime}^{\dagger}\mathbf{H}^{\prime(l)}, \mathbf{A}_{\Theta}^{\prime}^{\dagger} = \Psi(\mathbf{A}_{\Theta}^{\prime}), \mathbf{A}_{\Theta}^{\prime}^{\dagger}\mathbf{A}_{\Theta}^{\prime} \approx \mathbf{I},$$
(10)

where the goal is to learn an inverse operator Ψ that approximates $\mathbf{A}_{\Theta}'^{\dagger}$, and the regularization term $\mathbf{A}_{\Theta}'^{\dagger}\mathbf{A}_{\Theta}'\approx\mathbf{I}$ enforces an inverse constraint. We learn this operator because analytical inversion can be computationally expensive and unstable, especially as \mathbf{A}_{Θ}' may vary across samples and training iterations. A learned Ψ provides an approximation that generalizes across boundary structures. In

implementation, we can parameterize Ψ using a differentiable architecture such as a GNN or low-rank factorization. For the explicit inverse-operator learner, see Algorithm 8.7.3 in Appendix 8.7.

5 EXPERIMENTS

Datasets. We evaluate on two directed-network datasets. (i) River A real-world river network preprocessed from LamaH-CE2 (Klingler et al., 2021) over the Danube basin, providing hourly discharge and meteorological records. The graph has 358 nodes and 357 directed edges, partitioned into 209 boundary and 149 interior nodes. Each node has five features: discharge, surface air pressure, precipitation, temperature, and soil moisture. (ii) Blood flow A simulated arterial-network dataset generated with openBF (Benemerito et al., 2024), on a Circle of Willis model (Vrselja et al., 2014); we use a connected subnetwork with 14 nodes and 14 directed edges. Each node carries four features: flux, pressure, velocity, and cross-sectional area. For both datasets, model inputs are formed by concatenating the past W hours of features along the channel dimension; the prediction target is flux $\mathbf{y} \in \mathbb{R}^{|V|}$ at horizon t+n. All variables are independently normalized via z-score to ensure consistent scaling across nodes and variables.

Metrics. We evaluate models under a supervised node regression setup, following (Kirschstein & Sun, 2024; Jiang et al., 2025). Given W hours of historical flux data for all nodes, the goal is to predict the flux volume n hours ahead. In our setting, we use W=24 and a forecast horizon of n=6 hours. We compute the mean-squared error (MSE) on three node sets: all nodes $\ell(\hat{\mathbf{y}},\mathbf{y})=\frac{1}{|V|}\sum_{v_i\in V}(\hat{y}_i-y_i)^2$, boundary nodes $\ell_{BN}(\hat{\mathbf{y}},\mathbf{y})=\frac{1}{|V_{\rm BN}|}\sum_{v_i\in V_{\rm BN}}(\hat{y}_i-y_i)^2$, and interior nodes $\ell_{IN}(\hat{\mathbf{y}},\mathbf{y})=\frac{1}{|V_{\rm IN}|}\sum_{v_i\in V_{\rm IN}}(\hat{y}_i-y_i)^2$.

Competitors. We compare with several GNN baselines, including residual variants of graph convolutional networks (ResGCN) (Wu et al., 2019), graph attention networks (ResGAT) (Veličković et al., 2017) and GCNII (Chen et al., 2020). They also serve as backbone for ghost-TFP. We further compare against a dense graph transformation as ablation study (Wang et al., 2025).

Implementation. We implement the Robin-type boundary condition to model the structural coupling among v_g , v_b , and $v_{\rm nbr}$, as defined in Eq. (5). The message-passing layers are implemented using ResGCN, ResGAT and GCNII, with 128 dimensional node embedding, applying ReLU activation to every non-linear layer. We adapt the Picard search method (Paniconi & Putti, 1994) to accelerate the implicit GNN training, following Gu et al. (2020). We use a fully connected and trainable adjacency matrix to learn the inverse operator. For ablation study, we learn it over symmetric (i.e., $\mathbf{A}'_{\Theta}^{\dagger}$ in Eq. (10)) and asymmetric (i.e., $\mathbf{A}'^{\dagger}\Theta$, \mathbf{A}' in Eq. (6) and Θ in Eq. (9)) versions, following Wang et al. (2025), with detailed results and analysis deferred to RQ5. For the compared models, we implement their GNN architectures following Kirschstein & Sun (2024) and evaluate their results on both ground-truth, forward flow (i.e., \mathbf{A}) and the physically-implausible, reverse flow (i.e., \mathbf{A}^{\top}), which enables to verify whether adding boundary condition will improve their topology-awareness during training, with results deferred to RQ2.

RESULTS AND ANALYSIS

Based on the results in Table 2, 3 and 4, we answer the following research questions (**RQ1-5**):

RQ1. To what extent can ghost nodes compensate boundary information deficit in GNN training?

We answer this question by analyzing the improvement of adding ghost nodes in existing GNN backbones even in naive ways. Specifically, we implement Eq. (5) to use an MLP to learn and add ghost nodes, link them to their corresponding boundary nodes, and then run standard message passing. For River, before adding ghost nodes, standard GNNs often underperform the topology-agnostic MLP due to high errors on boundary nodes. Adding ghost nodes directly reduces boundary MSE across all backbones by 10.7% for ResGAT, 6.7% for ResGCN, and 15.7% for GCNII, and cuts overall Test MSE by 8.83% for ResGAT, 5.68% for ResGCN, and 10.53% for GCNII, bringing them below the MLP baseline (0.1135). For Blood, adding ghost nodes reduces Boundary MSE by 7.01% / 5.96% / 7.99% for ResGAT / ResGCN / GCNII and reduces Avg MSE by 5.19% / 4.57% / 5.34%. All ghost-enhanced backbones are below the MLP baseline (0.0670). These results confirm that by correctly modeling boundaries, ghost nodes consistently compensate for boundary-information deficits and allow GNNs to outperform topology-agnostic models across both datasets.

RQ2. Will adding ghost nodes improve the topology-awareness of standard GNNs?

Table 2: MSE comparison for River (left) and Blood (right), with Boundary/Interior breakdown. Shaded rows show group means: Avg (Base) averages the three baselines (GCNII, ResGCN, ResGAT), and $ghost-TFP_{Avg}$ averages their ghost counterparts. Diff (%) is the relative difference between boundary and interior.

Flux Predictors	River				Blood			
riux i redictors	Avg.	Boundary	Interior	Diff. (%)	Avg.	Boundary	Interior	Diff. (%)
GCNII	0.1253	0.1511	0.0891	41.03	0.0674	0.1451	0.0363	74.98
ResGCN	0.1232	0.1408	0.0985	30.04	0.0569	0.1140	0.0341	70.09
ResGAT	0.1166	0.1401	0.0836	40.33	0.0482	0.1056	0.0252	76.14
Avg (Base)	0.1217	0.1440	0.0904	37.13	0.0575	0.1216	0.0319	73.74
ghost-TFP _{GCNII}	0.1121	0.1274	0.0906	28.89	0.0638	0.1335	0.0359	73.11
ghost-TFP _{ResGCN}	0.1162	0.1313	0.0950	27.65	0.0543	0.1072	0.0331	69.12
ghost-TFP _{ResGAT}	0.1063	0.1251	0.0800	36.05	0.0457	0.0982	0.0247	74.85
ghost-TFP _{Avg}	0.1115	0.1279	0.0885	30.86	0.0546	0.1130	0.0312	72.36

Table 3: Topological comparison of Fwd, Rev, and ghost-TFP on River and Blood. In each cell, the first line shows Avg with its change relative to Rev; the second line (with results parenthesized) shows Boundary with its change relative to Rev. ↑ indicates better (lower MSE), and ↓ indicates worse (higher MSE).

		River			Blood	
Backbone	Fwd	Rev	ghost-TFP	Fwd	Rev	ghost-TFP
GCNII	0.1253 († 2.8%)	0.1289	0.1121 († 13.0%)	0.0674 (\psi 3.7%)	0.0650	0.0638 († 1.8%)
GCMII	$(0.1511 (\downarrow 1.5\%))$	(0.1488)	$(0.1274 (\uparrow 14.4\%))$	$(0.1451 (\downarrow 9.1\%))$	(0.1330)	$(0.1335 (\downarrow 0.4\%))$
ResGCN	0.1232 († 1.0%)	0.1245	$0.1162 (\uparrow 6.7\%)$	0.0569 (\psi 3.1%)	0.0552	0.0543 († 1.6%)
RESUCIA	$(0.1408 (\downarrow 2.8\%))$	(0.1369)	$(0.1313 (\uparrow 4.1\%))$	$(0.1140 (\downarrow 9.4\%))$	(0.1042)	$(0.1072 (\downarrow 2.9\%))$
ResGAT	0.1166 († 1.1%)	0.1179	0.1063 († 9.8%)	0.0482 (\psi 4.6%)	0.0461	$0.0457 (\uparrow 0.9\%)$
RUSUAI	$(0.1401 (\downarrow 3.8\%))$	(0.1350)	$(0.1251 (\uparrow 7.3\%))$	$(0.1056 (\downarrow 9.9\%))$	(0.0961)	$(0.0982 (\downarrow 2.2\%))$

Yes. From Table 3, the ResGAT scores 0.1166 (Fwd) vs. 0.1179 (Rev), a tiny 0.0013 gap ($\approx 1.1\%$); after adding ghost nodes the forward MSE drops to 0.1063 while the reverse stays at 0.1179, widening the gap to 0.0116 (10.9%), i.e. $8.9 \times$. ResGCN's gap grows from 0.0013 (1.1%) to 0.0083 (6.7%), a $6.3 \times$ boost, and GCNII's from 0.0036 (2.9%) to 0.0168 (15.0%), a $4.7 \times$ boost. On Blood, the baseline shows a mild reverse-over-forward advantage: Fwd vs. Rev = +4.56% / +3.08% / +3.69% for ResGAT / ResGCN / GCNII ("+" means forward is worse). After adding ghost nodes, the forward models outperform the baseline reverse: Fwd vs. Rev = -0.87% / -1.63% / -1.85%, respectively. This indicates that ghost nodes remove boundary-induced confounding and restore direction-consistent behavior across datasets.

RQ3. How effective can the implicit GNN training as defined in Eq. (9) learn ghost nodes over naïve layer-wise message-passing?

Table 4 shows that the fixed-point training in Eq. (9) already improves over naïve propagation: on River, relative to the naïve baselines, it reduces the average error by 0.0080 and 0.0039 (-6.9% and -3.5%) compared with ghost-TFP_{ResGCN} and ghost-TFP_{ResGAT} remains roughly unchanged; on Blood, the implicit variant (w/ Ghost) achieves 0.0557, outperforming the naïve

Table 4: Ablation results on Boundary Nodes (BN) in MSE. White and gray rows indicate results from using the directional $({\bf A}'^{\dagger})$ or bidirectional $({\bf A}'_{\Theta}^{\dagger})$ inverse operators, respectively.

Inv. Op.	Ri	ver	Blood		
	w/o Ghost	w/ Ghost	w/o Ghost	w/ Ghost	
Implicit GNN	0.1119 (0.1273)	0.1082 (0.1236)	0.0586 (0.0743	3) 0.0557 (0.0708)	
ResGAT + $\mathbf{A'}^{\dagger}$ ResGAT + $\mathbf{A'_{\Theta}}^{\dagger}$				2) 0.0440 (0.0615) 7) 0.0434 (0.0608)	
ResGCN + $\mathbf{A'}^{\dagger}$ ResGCN + $\mathbf{A'}_{\Theta}^{\dagger}$				2) 0.0524 (0.0716) 3) 0.0516 (0.0708)	
GCNII + $\mathbf{A'}^{\dagger}$ GCNII + $\mathbf{A'}_{\Theta}^{\dagger}$				9) 0.0619 (0.0838) 1) 0.0612 (0.0825)	

ghost-TFP_{GCNII} (0.0638, -12.7%). Moving to the stronger inverse-operator learner with richer connectivity ${\bf A'_{\Theta}}^{\dagger}$ (Table 4) further reduces errors on both datasets: for River, the results are 0.1033 (ghost-TFP_{ResGAT}, -2.8%), 0.1127 (ghost-TFP_{ResGCN}, -3.0%), and 0.1040 (ghost-TFP_{GCNII}, -7.2%); for Blood, the results are 0.0434 (ghost-TFP_{ResGAT}, -5.0%), 0.0516 (ghost-TFP_{ResGCN}, -5.0%), and 0.0612 (ghost-TFP_{GCNII}, -4.1%). The two tables demonstrate that implicit fixed-point solvers exploit ghost-node information more effectively than naïve layer-wise propagation and keep improving with richer connectivity patterns ${\bf A'_{\Theta}}$.

RQ4. Can the explicit inverse operator learning reduce the runtime overhead of implicit computation without compromising prediction accuracy?

Yes. From Table 4, the explicit inverse-operator learner (A'_{Θ}) reduces ResGAT's MSE from 0.1082 (Implicit GNN w/ Ghost) to 0.1033 (4.5%) while shrinking the runtime ratio from 13.8 to 6.0, i.e., about $2.3 \times$ faster; on the GCNII backbone it lowers the error from 0.1082 to 0.1040 (3.9%) with a similar speed-up. On Blood, explicit inverse-operator learner maintains or improves accuracy, achieving relative improvements of 22.0% and 7.4% on ResGAT and ResGCN, respectively, over the Implicit GNN w/ Ghost reference (0.0557). The runtime acceleration is similar to that observed on River. These results indicate that directly learning the inverse operator achieves higher accuracy with significantly lighter computation than iteratively solving fixed-point equations (Gu et al., 2020).

RQ5. Is it better to learn inverse over the parametric adjacency A'_{Θ} or its non-parametric counterpart A', why?

Table 4 shows that, with Ghost nodes enabled, replacing the symmetric adjacency \mathbf{A}_{Θ}' with its directed counterpart \mathbf{A}' consistently reduces the MSE: ResGAT drops from 0.1057 to 0.1033 (2.3%), ResGCN from 0.1151 to 0.1127 (2.1%), and GCNII from 0.1109 to 0.1040 (6.2%). On Blood, the same preference holds with smaller margins: ResGAT drops from 0.0440 to 0.0434 (1.36%), ResGCN from 0.0524 to 0.0516 (1.53%), and GCNII from 0.0619 to 0.0612 (1.13%).

Although the physical flow graph is strictly downstream-oriented, our bidirectional dense formulation better matches the *global* fixed-point operator implicit GNNs approximate, enabling long-range couplings beyond the observed sparse topology. This agrees with the benefits of dense graph transformations reported by Wang et al. (2025). Physically, both domains are open systems with unobserved exchanges (e.g., rainfall/groundwater/withdrawals in rivers; collateral and micro-circulatory paths in vasculature). Allowing bidirectional edges in the learned adjacency helps absorb these latent inflow—outflow processes, relaxes overly strict conservation biases in the observed graphs, and improves predictive accuracy across datasets.

6 RELATED WORK

Graph Augmentation with Virtual Nodes. Adding a global virtual node to a graph is a known technique to improve expressivity (Xu et al., 2019; Ying et al., 2021). This approach has been used to enhance graph-level prediction (Baek et al., 2021), reduce oversquashing (Hwang et al., 2022), and assist in physical simulations (Bentivoglio et al., 2025a; Mayr et al., 2023). However, a single global node does not address the localized boundary information deficit critical to our problem.

Implicit Graph Neural Networks (IGNNs). IGNNs compute node embeddings as a fixed point of a nonlinear system (Gu et al., 2020), a concept extended by deep equilibrium models (Bai et al., 2021; 2020). The IGNN paradigm has been analyzed through the lens of numerical diffusion (Chamberlain et al., 2021) and monotone operator theory (Baker et al., 2023), with various strategies proposed to mitigate issues like oversmoothing (Rusch et al., 2023). A key limitation remains their reliance on expensive, iterative solvers.

Boundary conditions in PIML. BCs are central in PIML: prior work balances PDE and BC losses via adaptive weights or architectural constraints and analyzes failure modes (Raissi et al., 2019; McClenny & Braga-Neto, 2020; Wang et al., 2023; Krishnapriyan et al., 2021). Operator-learning methods such as FNO often assume periodic BCs, whereas graph-based models better accommodate complex BCs (Li et al., 2021; Horie & Mitsume, 2022; Li et al., 2024). Other studies infer unknown BCs directly from data (Horuz et al., 2022; Zhao et al., 2022; Frerix et al., 2021). In contrast to approaches that fix BCs, our ghost-node formulation jointly learns boundary terms and interior dynamics in a data-driven manner (Liu et al., 2025).

7 CONCLUSION

This paper revisits the empirical shortcomings of GNNs in topological flux prediction and challenges the prevailing conclusion that GNNs are fundamentally unsuitable for such tasks. Through a surgery of the prediction loss behavior on fluid network, we demonstrate that the dominant source of error lies at boundary nodes. To compensate the boundary information deficit in GNN-based flux prediction, we propose a novel <code>ghost-TFP</code> framework, which augments GNNs with ghost nodes and an implicit solver to incorporate physically consistent boundary conditions during training. To improve scalability, we devise an explicit solver that learns inverse operators, enabling efficient layer-wise computation. Experiment demonstrates that <code>ghost-TFP</code> improves predictive accuracy and reduces the boundary-interior loss gap across multiple standard GNN backbones.

REFERENCES

- Jinheon Baek, Minki Kang, and Sung Ju Hwang. Accurate learning of graph representations with graph multiset pooling. In *International Conference on Learning Representations*, 2021.
- Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Multiscale deep equilibrium models. In *Advances in Neural Information Processing Systems*, 2020.
 - Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Stabilizing equilibrium models by jacobian regularization. In *International Conference on Machine Learning*, 2021.
 - Justin Baker, Qingsong Wang, Cory D Hauck, and Bao Wang. Implicit graph neural networks: A monotone operator viewpoint. In *International Conference on Machine Learning*, pp. 1521–1548. PMLR, 2023.
 - I Benemerito, A Melis, Antoine Wehenkel, and A Marzo. openbf: an open-source finite volume 1d blood flow solver. *Physiological Measurement*, 45(12):125002, 2024.
 - Roberto Bentivoglio, Elvin Isufi, Sebastiaan N. Jonkman, and Riccardo Taormina. Multi-scale hydraulic graph neural networks for flood modelling. *Natural Hazards and Earth System Sciences*, 25:335–351, 2025a.
 - Roberto Bentivoglio, Elvin Isufi, Sebastiaan Nicolas Jonkman, and Riccardo Taormina. Multiscale hydraulic graph neural networks for flood modelling. *Natural Hazards and Earth System Sciences*, 25(1):335–351, 2025b.
 - Vasile Berinde. *Iterative approximation of fixed points*. Springer, 2007.
 - Abraham Berman and Robert J Plemmons. *Nonnegative matrices in the mathematical sciences*. SIAM, 1994.
 - Christian Busse, Andrew P Kach, and Stephan M Wagner. Boundary conditions: What they are, how to explore them, why we need them, and when to consider them. *Organizational Research Methods*, 20(4):574–609, 2017.
 - Benjamin P. Chamberlain, Nick Rowbottom, James Li, and Michael M. Bronstein. Grand: Graph neural diffusion. In *International Conference on Machine Learning (ICML)*, pp. 1407–1419, 2021.
 - Kaixuan Chen, Shunyu Liu, Tongtian Zhu, Ji Qiao, Yun Su, Yingjie Tian, Tongya Zheng, Haofei Zhang, Zunlei Feng, Jingwen Ye, et al. Improving expressivity of gnns with subgraph-specific factor embedded normalization. In *KDD*, pp. 237–249, 2023.
 - Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. Simple and deep graph convolutional networks. In *International conference on machine learning*, pp. 1725–1735. PMLR, 2020.
 - David P Chock. A comparison of numerical methods for solving the advection equation—iii. *Atmospheric Environment. Part A. General Topics*, 25(5-6):853–871, 1991.
 - Thomas Frerix, Dmitrii Kochkov, Jamie A Smith, Daniel Cremers, Michael P Brenner, and Stephan Hoyer. Variational data assimilation with a learned inverse observation operator. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2021.
 - W Rockwell Geyer and Parker MacCready. The estuarine circulation. *Annual review of fluid mechanics*, 46(1):175–197, 2014.
 - Fangda Gu, Heng Chang, Wenwu Zhu, Somayeh Sojoudi, and Laurent El Ghaoui. Implicit graph neural networks. In *Advances in neural information processing systems*, volume 33, pp. 11984–11995, 2020.
 - Masanobu Horie and Naoto Mitsume. Physics-embedded neural networks: Graph neural pde solvers with mixed boundary conditions. In *Advances in Neural Information Processing Systems*, 2022. NeurIPS.

- Coşku Can Horuz, Matthias Karlbauer, Timothy Praditia, Martin V Butz, Sergey Oladyshkin, Wolfgang Nowak, and Sebastian Otte. Infering boundary conditions in finite volume neural networks.
 In International Conference on Artificial Neural Networks (ICANN), 2022.
 - EunJeong Hwang, Veronika Thost, Shib Sankar Dasgupta, and Tengfei Ma. An analysis of virtual nodes in graph neural networks for link prediction. In *The first learning on graphs conference*, 2022.
 - Haoyang Jiang, Jindong Wang, Xingquan Zhu, and Yi He. Topology-aware neural flux prediction guided by physics. *ICML*, 2025.
 - Guangyin Jin, Yuxuan Liang, Yuchen Fang, Zezhi Shao, Jincai Huang, Junbo Zhang, and Yu Zheng. Spatio-temporal graph neural networks for predictive learning in urban computing: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 2023.
 - Olaf Jonkeren, Piet Rietveld, and Jos van Ommeren. Climate change and inland waterway transport: welfare effects of low water levels on the river rhine. *Journal of Transport Economics and Policy (JTEP)*, 41(3):387–411, 2007.
 - Nikolas Kirschstein and Yixuan Sun. The merit of river network topology for neural flood forecasting. In *ICML*, 2024.
 - Christoph Klingler, Karsten Schulz, and Mathew Herrnegger. Lamah— large-sample data for hydrology and environmental sciences for central europe. *Earth System Science Data Discussions*, 2021:1–46, 2021.
 - Frederik Kratzert, Daniel Klotz, Martin Gauch, Christoph Klingler, Grey Nearing, and Sepp Hochreiter. Large-scale river network modeling using graph neural networks. In *EGU General Assembly Conference Abstracts*, pp. EGU21–13375, 2021.
 - Aditi S. Krishnapriyan, Amir Gholami, Shandian Zhe, Robert M. Kirby, and Michael W. Mahoney. Characterizing possible failure modes in physics-informed neural networks. In *Advances in Neural Information Processing Systems*, 2021.
 - Bernhard Lehner, Gregor Czisch, and Sara Vassolo. The impact of global change on the hydropower potential of europe: a model-based analysis. *Energy policy*, 33(7):839–855, 2005.
 - Randall J LeVeque. Finite difference methods for ordinary and partial differential equations: steady-state and time-dependent problems. SIAM, 2007.
 - Tianyu Li, Shufan Zou, Xinghua Chang, Xiaogang Deng, et al. Finite volume graph network (fvgn): Predicting unsteady incompressible fluid dynamics via boundary-aware message aggregation. *Physics of Fluids*, 2024.
 - Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew L. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.
 - Zewen Liu, Xiaoda Wang, Bohan Wang, Zijie Huang, Carl Yang, and Wei Jin. Graph odes and beyond: A comprehensive survey on integrating differential equations with graph neural networks. *arXiv preprint arXiv:2503.23167*, 2025.
 - Donald C Mangold and Chin-Fu Tsang. A summary of subsurface hydrological and hydrochemical models. *Reviews of Geophysics*, 29(1):51–79, 1991.
 - Andreas Mayr, Sebastian Lehner, Arno Mayrhofer, Christoph Kloss, Sepp Hochreiter, and Johannes Brandstetter. Boundary graph neural networks for 3d simulations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 9099–9107, 2023.
 - Liam McClenny and Ulisses Braga-Neto. Self-adaptive physics-informed neural networks using a soft attention mechanism. In *AAAISpring Symposium on Combining Learning and Reasoning*, 2020.

- Claudio Paniconi and Mario Putti. A comparison of picard and newton iteration in the numerical solution of multidimensional variably saturated flow problems. *Water Resources Research*, 30 (12):3357–3374, 1994.
 - N LeRoy Poff, J David Allan, Mark B Bain, James R Karr, Karen L Prestegaard, Brian D Richter, Richard E Sparks, and Julie C Stromberg. The natural flow regime. *BioScience*, 47(11):769–784, 1997.
 - Michael Poli, Stefano Massaroli, Junyoung Park, Atsushi Yamashita, Hajime Asama, and Jinkyoo Park. Graph neural ordinary differential equations. *arXiv preprint arXiv:1911.07532*, 2019.
 - K Manjunatha Prasad and RB Bapat. The generalized moore-penrose inverse. *Linear Algebra and its Applications*, 165:59–69, 1992.
 - Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear pdes. *Journal of Computational Physics*, 378:686–707, 2019.
 - Jie Ren, Xidong Feng, Bo Liu, Xuehai Pan, Yao Fu, Luo Mai, and Yaodong Yang. Torchopt: An efficient library for differentiable optimization. *Journal of Machine Learning Research*, 24(367): 1–14, 2023.
 - T. Konstantin Rusch, Michael M. Bronstein, and Siddhartha Mishra. A survey on oversmoothing in graph neural networks. *arXiv preprint arXiv:2303.10993*, 2023.
 - Kathrin Schacke. On the kronecker product. Master's thesis, University of Waterloo, 2004.
 - Yu-Heng Tseng and Joel H Ferziger. A ghost-cell immersed boundary method for flow in complex geometry. *Journal of computational physics*, 192(2):593–623, 2003.
 - Anna Varbella, Kenza Amara, Blazhe Gjorgiev, Mennatallah El-Assady, and Giovanni Sansavini. Powergraph: A power grid benchmark dataset for graph neural networks. *Advances in Neural Information Processing Systems*, 37:110784–110804, 2024.
 - Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
 - Zvonimir Vrselja, Hrvoje Brkic, Stefan Mrdenovic, Radivoje Radic, and Goran Curic. Function of circle of willis. *Journal of Cerebral Blood Flow & Metabolism*, 34(4):578–584, 2014.
 - Hongjun Wang, Jiyuan Chen, Yinqiang Zheng, and Xuan Song. Accelerating flood warnings by 10 hours: the power of river network topology in ai-enhanced flood forecasting. *npj Natural Hazards*, 2(1):45, 2025.
 - Jian Wang, Yifan Mo, Bashar Izzuddin, and Chang-Won Kim. Exact dirichlet boundary physics-informed neural network (epinn) for solid mechanics. *Computer Methods in Applied Mechanics and Engineering*, 414:116184, 2023.
 - Paul G Whitehead, Robert L Wilby, Richard W Battarbee, Martin Kernan, and Andrew John Wade. A review of the potential impacts of climate change on surface water quality. *Hydrological sciences journal*, 54(1):101–123, 2009.
- Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *International conference on machine learning*, pp. 6861–6871. PMLR, 2019.
- Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S Yu. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020.
 - Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations (ICLR)*, 2019.

Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform badly for graph representation? *Advances in neural information processing systems*, 34:28877–28888, 2021.

Qingqing Zhao, David B Lindell, and Gordon Wetzstein. Learning to solve pde-constrained inverse problems with graph networks. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, 2022.

Baodong Zheng and Liancheng Wang. Spectral radius and infinity norm of matrices. *Journal of mathematical analysis and applications*, 346(1):243–250, 2008.

Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. AI open, 1:57–81, 2020.

8 APPENDIX

8.1 ETHICS STATEMENT

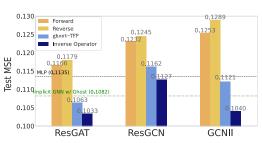
We adhere to the ICLR Code of Ethics. Our study uses only publicly available datasets or software, and involves no human subjects or sensitive attributes beyond what is already released. We audited potential risks (misuse, bias, privacy) and found no foreseeable harms specific to our methods; dataset licenses and legal compliance were respected. There are no conflicts of interest or external sponsorship that influenced this work.

8.2 Reproducibility Statement

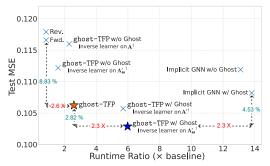
We aim for full reproducibility. The paper specifies model architectures and evaluation protocols; ablation settings are detailed in Section 5. All datasets and preprocessing steps are documented in Section 5. We provide anonymized code with fixed random seeds in the supplementary materials to enable exact replication of results.

8.3 Supplementary for Experiments

These figures complement the main results by contrasting the message flow topology and efficiency. The bar chart compares Forward and Reverse variants across backbones, showing small gaps between the two, while <code>ghost-TFP</code> and the inverse learner consistently reduce test MSE. The scatter plot relates test MSE to runtime, where points closer to the lower-left indicate better accuracy-efficiency trade-offs; <code>ghost-TFP</code> occupies this region.



(a) Forward vs. Reverse GNNs (left two bars) highlight the similarity from different message flow direction, while ghost-TFP and our inverse learner (right two bars) drive errors lower and widen the gap with reverse flow.



(b) Each point shows model's test MSE against its runtime ratio relative to the ResGAT Forward baseline (1×). Points that fall lower and further left deliver both lower error and faster inference, highlighting superior accuracy—efficiency trade-offs.

8.4 From Closed Subgraphs to Open Subsystems

Problem overview. The 14 evaluation nodes originate from a larger, 30-node system. Training and evaluating them as an isolated subgraph implicitly imposes a closed-system inductive bias: only intra-graph interactions are modeled, while upstream/downstream exchanges and external forcings are omitted. This bias disproportionately affects Boundary nodes, where missing cross-boundary inputs manifest as inflated errors. To mitigate this, we introduce ghost nodes that learn a data-driven proxy for the absent external coupling at the boundary, effectively "half-opening" the subgraph.

Discussion. Table 5 shows that when the same 14 nodes are trained and evaluated as a closed subgraph, the boundary penalty relative to the overall average is very high (+82.6%), indicating that the omission of external interactions disproportionately affects boundary behavior. Evaluating the identical nodes inside the full graph substantially reduces this penalty to +10.6%, reflecting the richer upstream/downstream context and constraints available in the larger system. When only the

Table 5: Results for the same 14 nodes under three settings (Subgraph / FullGraph / Ghost-Subgraph). $\%\Delta_{Boundary}$ denotes the boundary increase relative to Avg, with smaller values indicating a weaker boundary penalty.

Setting	Avg	Interior	Boundary	$\% \Delta_{ ext{Boundary}}$
Subgraph	0.0530	0.0355	0.0968	+82.6%
Fullgraph	0.0930	0.0891	0.1029	+10.6%
Ghost-Subgraph	0.0476	0.0342	0.0811	+70.4%

subgraph is available, adding ghost nodes still narrows the gap by learning a data-driven boundary proxy: the penalty drops from +82.6% to +70.4%, while the overall Avg also improves. Altogether, these results support the view that closed subgraphs amplify boundary difficulty, embedding the subsystem in the larger graph naturally balances errors, and ghost nodes provide a lightweight way to approximate external interactions when access to external nodes is not feasible.

8.5 Derivation for $\mathbf{A}_{\mathbf{\Theta}}'$

Notation (Symbol Table). To keep the appendix consistent with the main text, we list the symbols used below. Only the Robin coefficients have been renamed from (α, β) to (w_1, w_2) ; the advection speed a is unchanged, and the spatial step is uniformly denoted by Δx .

- u(x,t): scalar state; u_i^n is the discrete state at grid index i and time level n.
- a > 0: advection speed in the governing PDE (kept as is).
- Δx : spatial grid spacing; Δt : time step.
- $\sigma := \frac{a \Delta t}{\Delta x}$: Courant number for advection.
- w_1, w_2 : Robin boundary coefficients (replacing α, β) used only in the boundary condition.
- u_L : prescribed boundary trace entering the Robin condition; when needed we use u_L^{n+1} to indicate the time level.
- A'_Θ: implicit system matrix assembled from interior and boundary discrete equations at time level n + 1.

This section provides a detailed derivation for the components of the implicit system matrix A'_{Θ} , establishing the theoretical foundation for the ghost-TFP framework discussed in the main paper. By using the 1D advection equation as a canonical example, this derivation serves to:

- Justify the interpretation of GNN message-passing as a numerical discretization of a physical system's spatial dynamics.
- Demonstrate how boundary conditions introduce specific mathematical constraints that are absent in standard GNN formulations.
- Show how these discrete equations naturally form the A'_{Θ} , which is the core problem our implicit solver addresses.

The derivations are based on two fundamental principles:

• The Governing PDE: The 1D advection equation, $\frac{\partial u}{\partial t} + a \frac{\partial u}{\partial x} = 0$, where a > 0.

• The Boundary Condition: The Robin-type condition at the boundary x=0, given by $w_1u+w_2\frac{\partial u}{\partial x}=u_L$.

8.5.1 Derivation of the Interior Equations

Here is the derivation of the interior equations:

Target Equation:

$$(1+\sigma)u_i^{n+1} - \sigma u_{i-1}^{n+1} = u_i^n$$
 for $i = 1, \dots, N$

Method: Apply the standard implicit first-order upwind scheme to the governing PDE at an interior node i.

1. **Discretize the Time Derivative:** We approximate the partial derivative with respect to time, $\frac{\partial u}{\partial t}$, using a first-order forward difference:

$$\frac{\partial u}{\partial t} \approx \frac{u_i^{n+1} - u_i^n}{\Delta t}$$

2. Discretize the Spatial Derivative (Implicitly): The term "implicit" signifies that the spatial derivative is evaluated at the future time step, n+1. The term "upwind" (for a>0) means we use a backward difference, looking at the node from which the flow originates (i-1).

$$\frac{\partial u}{\partial x} \approx \frac{u_i^{n+1} - u_{i-1}^{n+1}}{\Delta x}$$

3. **Combine and Simplify:** Substitute these approximations back into the governing PDE, $\frac{\partial u}{\partial t} + a \frac{\partial u}{\partial x} = 0$:

$$\frac{u_i^{n+1} - u_i^n}{\Delta t} + a \left(\frac{u_i^{n+1} - u_{i-1}^{n+1}}{\Delta x} \right) = 0$$

Multiply the entire equation by Δt to clear the denominator:

$$(u_i^{n+1} - u_i^n) + \frac{a\Delta t}{\Delta x}(u_i^{n+1} - u_{i-1}^{n+1}) = 0$$

Let $\sigma = \frac{a\Delta t}{\Delta x}$ be the Courant number. Substituting σ gives:

$$u_i^{n+1} - u_i^n + \sigma(u_i^{n+1} - u_{i-1}^{n+1}) = 0$$

4. **Rearrange:** Group all the unknown terms (at time n + 1) on the left side and the known terms (at time n) on the right side.

$$u_i^{n+1} + \sigma u_i^{n+1} - \sigma u_{i-1}^{n+1} = u_i^n$$

Factoring out u_i^{n+1} yields the final target equation:

$$(1+\sigma)u_i^{n+1} - \sigma u_{i-1}^{n+1} = u_i^n$$

8.5.2 Derivation of the Boundary Condition Equation

Here is the Derivation of the Boundary Condition Equation:

Target Equation:

$$\left(w_1 - \frac{w_2}{\Delta x}\right)u_0^{n+1} + \left(\frac{w_2}{\Delta x}\right)u_1^{n+1} - u_L^{n+1} = 0$$

Method: This equation arises directly from discretizing the Robin boundary condition itself, without involving the PDE.

1. State the Boundary Condition: The Robin condition at node i=0 and at the future time step n+1 is:

$$w_1 u_0^{n+1} + w_2 \left(\frac{\partial u}{\partial x} \right) \Big|_{i=0}^{n+1} = u_L^{n+1}$$

2. **Discretize the Spatial Derivative:** At the boundary i = 0, we cannot use a backward difference. The natural choice is a first-order forward difference, using nodes 0 and 1:

$$\left(\frac{\partial u}{\partial x}\right)\Big|_{i=0}^{n+1} \approx \frac{u_1^{n+1} - u_0^{n+1}}{\Delta x}$$

(Here, we use Δx for the spatial step, consistent with the paper's notation).

3. **Combine and Rearrange:** Substitute the discretized derivative back into the boundary condition equation:

$$w_1 u_0^{n+1} + w_2 \left(\frac{u_1^{n+1} - u_0^{n+1}}{\Delta x} \right) = u_L^{n+1}$$

Distribute the term $w_2/\Delta x$:

$$w_1 u_0^{n+1} + \frac{w_2}{\Delta x} u_1^{n+1} - \frac{w_2}{\Delta x} u_0^{n+1} = u_L^{n+1}$$

Group the coefficients for u_0^{n+1} and move all terms to the left side to obtain the final form:

$$\left(w_1 - \frac{w_2}{\Delta x}\right)u_0^{n+1} + \left(\frac{w_2}{\Delta x}\right)u_1^{n+1} - u_L^{n+1} = 0$$

8.5.3 Derivation of the PDE at the Boundary

Here is the Derivation of the PDE at the Boundary:

Target Equation:

$$\left(1 - \frac{a\,\Delta t\,w_1}{w_2}\right)u_0^{n+1} + \left(\frac{a\,\Delta t}{w_2}\right)u_L^{n+1} = u_0^n$$

Method: This derivation cleverly combines the PDE and the BC. The key is to use the boundary condition to eliminate the spatial derivative term from the discretized PDE.

1. Discretize the PDE at the Boundary (i = 0): First, write the implicit discretization of the PDE at node i = 0:

$$\frac{u_0^{n+1} - u_0^n}{\Delta t} + a \left(\frac{\partial u}{\partial x}\right) \Big|_{i=0}^{n+1} = 0$$

This equation contains the spatial derivative term, which we need to handle.

2. **Isolate the Derivative from the Boundary Condition:** Return to the Robin condition from the previous section:

$$w_1 u_0^{n+1} + w_2 \left(\frac{\partial u}{\partial x}\right)\Big|_{i=0}^{n+1} = u_L^{n+1}.$$

We can rearrange this to solve for the derivative term:

$$w_{2} \left(\frac{\partial u}{\partial x} \right) \Big|_{i=0}^{n+1} = u_{L}^{n+1} - w_{1} u_{0}^{n+1}$$
$$\left(\frac{\partial u}{\partial x} \right) \Big|_{i=0}^{n+1} = \frac{u_{L}^{n+1} - w_{1} u_{0}^{n+1}}{w_{2}}$$

3. **Substitute and Simplify:** Now, substitute the expression for the derivative from Step 2 into the discretized PDE from Step 1:

$$\frac{u_0^{n+1} - u_0^n}{\Delta t} + a \left(\frac{u_L^{n+1} - w_1 u_0^{n+1}}{w_2} \right) = 0$$

Multiply the entire equation by Δt :

$$(u_0^{n+1} - u_0^n) + \frac{a\Delta t}{w_2}(u_L^{n+1} - w_1 u_0^{n+1}) = 0$$

Distribute the term $\frac{a\Delta t}{w_2}$:

$$u_0^{n+1} - u_0^n + \frac{a\Delta t}{w_2} u_L^{n+1} - \frac{a\Delta t \, w_1}{w_2} u_0^{n+1} = 0$$

4. **Rearrange:** Finally, group the unknown terms (n + 1) on the left side and the known term (n) on the right side.

$$\left(1 - \frac{a\Delta t \, w_1}{w_2}\right) u_0^{n+1} + \left(\frac{a\Delta t}{w_2}\right) u_L^{n+1} = u_0^n$$

In summary, the equations derived for the interior nodes (from the PDE) and the boundary nodes (from the boundary condition) collectively define the rows of the augmented system matrix \mathbf{A}_{Θ}' . The core contribution of the <code>ghost-TFP</code> framework lies in its ability to learn an efficient, implicit operator.

8.6 ILLUSTRATING THE AUGMENTED MATRIX

The core of our implicit solver revolves around constructing the system \mathbf{A}'_{Θ} . The structure of the system matrix \mathbf{A}'_{Θ} is a direct reflection of the underlying physical topology. This section provides concrete examples to illustrate this crucial connection. We will demonstrate the structure of \mathbf{A}'_{Θ} in complex, coupled ones.

8.6.1 MERGING THE SYSTEMS (7x7 COUPLED MATRIX)

Now, we alter the topology by introducing a new, seventh node, v_m , into which both systems merge. The flow paths become:

- $v_{b1} \rightarrow v_{d1} \rightarrow v_m$
- $v_{b2} \rightarrow v_{d2} \rightarrow v_m$

This introduces physical coupling, as the state of the merge point v_m now depends on inputs from both upstream branches. We augment our state vector with the new node:

$$\mathbf{H}'^{(\mathbf{t+1})} = \left[\begin{array}{c} \mathbf{h}_{g1}^{(t+1)}, \ \mathbf{h}_{b1}^{(t+1)}, \ \mathbf{h}_{d1}^{(t+1)}, \\ \mathbf{h}_{q2}^{(t+1)}, \ \mathbf{h}_{b2}^{(t+1)}, \ \mathbf{h}_{d2}^{(t+1)}, \ \mathbf{h}_{m}^{(t+1)} \end{array} \right]^{T}$$

The new 7×7 matrix, \mathbf{A}'_{Θ} , is no longer block-diagonal. The coupling appears precisely in the equation governing the new node v_m .

$$\mathbf{A}_{\Theta}' = \begin{pmatrix} 1 & -p_{1b} & -p_{1d} & 0 & 0 & 0 & 0 \\ -w_{b1,g1} & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & -w_{d1,b1} & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -p_{2b} & -p_{2d} & 0 \\ 0 & 0 & 0 & -w_{b2,g2} & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -w_{d2,b2} & 1 & 0 & 0 \\ 0 & 0 & -w_{m,d1} & 0 & 0 & -w_{m,d2} & 1 \end{pmatrix}$$

Analysis of the Coupling: The first six rows and columns largely retain the decoupled structure. The critical change is in the **last row**, which defines the update for v_m . The non-zero elements in columns 3 and 6, $-w_{m,d1}$ and $-w_{m,d2}$, are the mathematical signature of the physical merge. These elements, which reside in the previously zero off-diagonal block area, now link the two subsystems together through the dynamics of v_m . This clearly demonstrates how a local change in graph topology induces a specific, predictable change in the global system matrix.

8.7 ALGORITHMS

```
919
            Algorithm 8.7.1: Ghost-TFP with Standard GNN Backbone
920
                           : Graph G = (V, E) and its node features H, Set of boundary nodes V_{BN},
921
                             Number of GNN layers L, Upstream neighbor set of node v_i is \mathcal{N}(j),
922
                             Graph augmentation operator A, Ground-truth labels y,
923
                             Model components: GNN backbone f_{GNN}, Predictor MLP f_{pred}, Ghost MLP,
924
                             Loss function \mathcal{L} and Projection operator \Pi_A.
            Parameters: Learnable parameters \theta_{gb} for the Ghost MLP,
925
                             \theta_{GNN} for the GNN backbone, and \theta_{pred} for the predictor.
926
927
            V_{GH}, H_{GH}, E_{GH} \leftarrow \emptyset, \emptyset, \emptyset;
            for each boundary node v_b \in V_{BN} do
928
                 v_{nbr} \leftarrow v_{nbr} \in \{v_j \mid v_b \in \mathcal{N}(j)\};
                                                                                                    // Get interior neighbor
929
                 h_b, h_{nbr} \leftarrow H[v_b, :], H[v_{nbr}, :];
930
                 h_g \leftarrow \text{MLP}(\text{Concat}(h_b, h_{nbr}); \theta_{gb});
                                                                                                    // Learn ghost embedding
931
                 Create v_g and add to V_{GH}; Append h_g to H_{GH};
932
                 Add edge (v_g, v_b) to E_{GH}
933
            H' \leftarrow \text{Concat}(H, H_{GH});
                                                                                                         // Augmented features
            A' \leftarrow \mathcal{A}(V \cup V_{GH}, E \cup E_{GH});
934
                                                                                                        // Augmented adjacency
935
            H'^{(L)} \leftarrow f_{\text{GNN}}(A', H'; \theta_{\text{GNN}})
936
            H^{(L)} \leftarrow \Pi_A(H'^{(L)});
                                                                                   // Project back to original nodes
937
            \hat{y} \leftarrow f_{\text{pred}}(H^{(L)}; \theta_{\text{pred}})
938
            \mathcal{L}_{loss} \leftarrow \mathcal{L}(y, \hat{y});
939
            Update \theta_{gb}, \theta_{GNN}, \theta_{pred} using gradients of \mathcal{L}_{loss};
940
            return Prediction ŷ
941
942
943
944
            Algorithm 8.7.2: Implicit Ghost-Boundary Message-Passing
945
946
                           : Augmented graph G' = (A', H'), Ground truth y,
                             Predictor model f_{\text{pred}}, Loss function \mathcal{L}.
947
            Parameters: Shared weight matrix \Theta with constraint ||\Theta||_{\infty} \leq 1/\lambda_{pf}(A'),
948
                             Learnable parameters \theta_{pred} for the predictor.
949
                             Max iterations K, Convergence tolerance \epsilon > 0.
950
            H'^{(0)} \leftarrow H';
951
            \quad \text{for } k=1 \text{ to } K \text{ do}
952
                 H'^{(k)} \leftarrow \sigma(A'H'^{(k-1)}\Theta);
953
                 if ||H'^{(k)} - H'^{(k-1)}|| < \epsilon then
                  break;
955
            H'^* \leftarrow H'^{(k)}:
                                                            // Equilibrium embeddings satisfying Eq.
956
957
            H^* \leftarrow \Pi_A(H'^*);
                                                                             // Project back to original node set
            \hat{y} \leftarrow f_{\text{pred}}(H^*; \theta_{pred});
958
            \mathcal{L}_{loss} \leftarrow \mathcal{L}(y, \hat{y});
959
            \nabla_{H'^*} \mathcal{L}_{loss} \leftarrow \text{Compute gradient from loss up to the equilibrium point};
960
            q^{(0)} \leftarrow \mathbf{0}:
                                                                                       // Initialize implicit gradient
961
            D \leftarrow \sigma'(A'H'^*\Theta);
                                                                      // Jacobian of activation at equilibrium
962
            \quad \text{for } k=1 \text{ to } K \text{ do}
963
                 g^{(k)} \leftarrow D \odot ((A')^T g^{(k-1)} \Theta^T + \nabla_{H'^*} \mathcal{L}_{loss});
964
                 if ||g^{(k)} - g^{(k-1)}|| < \epsilon then
965
                  break;
966
967
                                                                                 // Converged implicit gradient 
abla_Z l
968
            Compute \nabla_{\Theta} \mathcal{L}_{loss} and \nabla_{\theta_{gh}} \mathcal{L}_{loss} using g^* via auto-differentiation;
969
            Update parameters \Theta, \theta_{gh}, \theta_{pred};
970
            return Prediction ŷ
```

```
978
979
980
981
982
983
984
985
986
987
988
989
            Algorithm 8.7.3: Explicit Adjacency-Inverse Solver
990
            Initialize: Augmented graph G' = (V', E'), Initial augmented features H',
991
                             Ground truth y, Loss function \mathcal{L}, Projection operator \Pi_A.
992
            Parameters: Number of layers L, Regularization strength \lambda,
993
                             Learnable parameters \theta_{\Psi} for Inverse operator \Psi ,
994
                             Parametric adjacency \theta_A for A_\Theta' , Predictor \theta_{pred} for f_{\mathrm{pred}} .
995
            Construct learnable parametric adjacency;
                                                                                          // upstream \leftrightarrow all downstream
996
            A'_{\Theta} using parameters \theta_A;
997
            (A'_{\Theta})^{\dagger} \leftarrow \Psi(A'_{\Theta}; \theta_{\Psi});
                                                              // Approximate inverse with learned operator
998
            H'^{(0)} \leftarrow H':
999
            for l=0 to L-1 do
1000
             H'^{(l+1)} \leftarrow \sigma((A'_{\Theta})^{\dagger}H'^{(l)});
                                                                           // Apply GNN layer with non-linearity
1001
            H^{(L)} \leftarrow \Pi_A(H'^{(L)});
1002
                                                                             // Project back to original node set
1003
            \hat{y} \leftarrow f_{\text{pred}}(H^{(L)}; \theta_{pred});
            \mathcal{L}_{\text{node}} \leftarrow \mathcal{L}(y, \hat{y});
                                                                                                      // Node prediction loss
            \mathcal{L}_{\text{reg}} \leftarrow \lambda || (A'_{\Theta})^{\dagger} A'_{\Theta} - I ||_F^2 ;
1005
                                                                             // Inverse constraint regularization
            \mathcal{L}_{total} \leftarrow \mathcal{L}_{node} + \mathcal{L}_{reg};
1006
1007
            Update \theta_{\Psi}, \theta_{A}, \theta_{pred} using gradients of \mathcal{L}_{total};
            return Prediction \hat{y}
1008
1009
1010
```