Physics-Informed Implicit Representations of Network Flows

Anonymous Author(s)
Affiliation
Address
email

Abstract

Flow networks are ubiquitous in natural and engineered systems, and in order to understand and manage these networks, one must quantify the flow of commodities across their edges. This paper considers the estimation problem of predicting unlabeled edge flows from nodal supply and demand. We propose an implicit neural network layer that incorporates two fundamental physical laws: conservation of mass, and the existence of a constitutive relationship between edge flows and nodal states (e.g., Ohm's law). Computing the edge flows from these two laws is a nonlinear inverse problem, which our layer solves efficiently with a specialized contraction mapping. Using implicit differentiation to compute the solution's gradients, our model is able to learn the constitutive relationship within a semi-supervised framework. We demonstrate that our approach can accurately predict edge flows in several experiments on AC power networks.

1 Introduction

Network flows are a fundamental aspect of modern society, from traffic and communication networks to power and water distribution systems. Many critical infrastructures are well-modeled as graphs, with edges that transport vital commodities [1]. Beyond infrastructure, network flows are also central to models in epidemiology, ecology, medicine, and chemical networks. Their dynamics have been well-studied in compartmental systems theory [2]. Given the prevalence of flow networks in natural and engineered systems, predicting flows in these networks is an important learning task that may facilitate monitoring, control, optimization, and protection of these networks.

While domain-specific tools to predict network flows have been around for a while, the machine learning community has only recently taken an interest in general-purpose models for network flows. [3] predicts edge flows from partial measurements by making a smoothing assumption, i.e., by minimizing nodal flow divergence. [4] improves on this approach by adding a trainable regularizer that can incorporate side information. Both of these approaches are centered on a notion of approximate conservation, i.e., that the net inflow to each node should be near zero. Since conservation of mass is a universal constraint on network flows, imposing this conservation law is an important step toward embedding physics into the model.

But the conservation law alone is not enough to uniquely determine flow, which is why both [3] and [4] rely on heuristic regularizers to select the "best" conservation-respecting flow. In fact, physical networks are often governed by a *pair* of physical laws: the conservation law, and a *constitutive relationship*, which specifies the magnitude and direction of each edge flow based on "effort" variables at each incident node (e.g., pressure or voltage). For example, in DC circuits, currents are conserved according to Kirchoff's current law, and Ohm's law is the constitutive relationship that relates current flows to nodal potentials. The conservation law and the constitutive relationship together define the unique edge flows (and nodal efforts).

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

1.1 Contributions

This paper proposes a model for network flows that embeds both the conservation law and existence of a constitutive relationship. Our model, which we call an *Implicit Flow Network* (IFN), predicts each edge flow using a trainable nonlinear function of latent nodal variables. These latent variables are constrained to a manifold wherein the conservation law is satisfied. In addition to introducing IFN, we offer the following contributions: (i) a contraction algorithm that is able to both evaluate the IFN layer and backpropagate gradients through it, (ii) an explicit upper bound on the number of iterations required by this algorithm, (iii) a rigorous theoretical comparison between IFN and the state-of-the-art flow estimation methods in [3, 4], and (iv) numerical experiments from several AC power networks that indicate IFN can significantly outperform these baselines on the flow estimation task. Additionally, because IFN requires a nonlinearity with a constrained slope, we provide (v) a novel "derivative-constrained perceptron", which is essentially a trainable activation function with upper and lower bounds on its slope.

1.2 Related Work

Network Flow Estimation Flows on graphs are a classical topic in computer science [5], and flow forecasting has long been studied in specific domains like traffic [6], but interest in the flow estimation task from a machine learning perspective appears to be relatively recent. Deep learning algorithms have been used to predict traffic flows [7, 8], but [3] and [4] appear to be the first papers to propose methods for domain-agnostic flow prediction, based on the notion of divergence minimization.

Implicit Neural Networks IFN belongs to a growing class of models called implicit neural networks, which do not explicitly state the output of the model; rather, they describe a desired relationship between the model's inputs and outputs. In the prevailing implicit framework, the output is defined as a fixed point of a trainable perceptron. This approach was introduced in [9] as a "deep equilibrium network". Subsequent work has developed new frameworks for ensuring the existence of the fixed point and computing it [10, 11, 12, 13, 14]. Other types of implicit neural networks include neural ODEs [15] and layers that solve convex optimization problems [16].

Graph Neural Networks Graph neural networks (GNN) are a diverse family of models for network-related learning tasks that incorporate graph structure directly into the model. GNNs can typically be classified into three types, in increasing order of generality [17, §5.3]: convolutional models [18, 19], attentional models [20], and message-passing models [21, 22]. Recently, [23] proposed an implicit graph convolutional network. Analogously, IFN can be interpreted as an implicit message-passing GNN, with flows serving as messages and latent nodal variables acting as an embedding.

1.3 Preliminaries and Notation

Given a directed graph $G = (\mathcal{V}, \mathcal{E})$, the signed incidence matrix $B \in \{-1, 0, 1\}^{|\mathcal{V}| \times |\mathcal{E}|}$ is the matrix with entries

$$B_{i,e} = \begin{cases} 1, & i \text{ is the head of } e \\ -1, & i \text{ is the tail of } e \\ 0, & \text{else} \end{cases}, \quad \forall i \in \mathcal{V} \text{ and } e \in \mathcal{E}$$

For an undirected graph, the signed incidence matrix is obtained by assigning an aribtrary orientation to each edge.

Given a vector $x \in \mathbb{R}^n$, we use the notation $[x]$ to denote the diagonal matrix $\mathrm{diag}(x) \in \mathbb{R}^{n \times n}$. Where such notation would be unclear (e.g., may be confused with brackets to indicate order of operations), we fall back on the $\mathrm{diag}(\cdot)$ notation. We write $x^\perp$ to refer to the vector space that is orthogonal to x, i.e., the space $\{x' \in \mathbb{R}^n : x^\top x' = 0\}$. Given a positive definite diagonal matrix $D \in \mathbb{R}^{n \times n}$, we write $||x||_{2,D}$ to represent the weighted 2-norm $||D^{\frac{1}{2}} x||_2$. Given any matrix M, we use the notation M_i to represent the ith column vector of M, and $M^{(j)}$ indicates the transpose of the jth row vector.

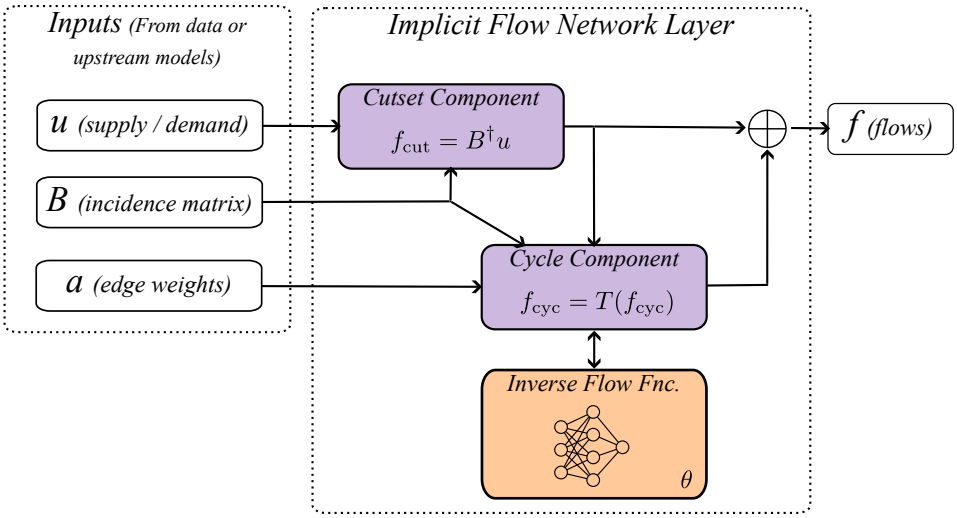

Figure 1: Diagram of the IFN. Inputs are the supply / demand vector u, incidence matrix B, and edges weights a, which are either known or output from upstream models. The IFN layer separately computes the cutset component and cycle component of the flows, with a trainable model for the inverse of the flow function in the constitutive relationship. These components are summed and output as the flow prediction, for downstream use.

2 Implicit Flow Networks

Implicit flow networks are a family of message-passing GNNs that are inspired by the physics of network systems. In many physical networks, nodes "communicate" through the exchange of a commodity, like power, water, or force. These exchanges of commodities can be represented as flows across edges in the graph. These flows tend to be conserved, particularly in the steady states of the system, so that Kirchoff's current law is satisfied at every node. Furthermore, these flows are determined by a constitutive equation that yields the magnitude and direction of an edge flow based on states of the incident nodes. Implicit flow networks replicate these two constraints.

Definition 2.1 (Implicit Flow Network)**.** An *implicit flow network* (IFN) is a module with the following components:

1. fixed parameters $0 < d_{\min} \leq d_{\max}$,
2. trainable parameters $\theta \in \mathbb{R}^r$ for some r, and
3. a family of differentiable functions $h_\theta : \mathbb{R} \to \mathbb{R}$ such that $d_{\min} \leq h'_\theta(y) \leq d_{\max}$ for all $y \in \mathbb{R}$ and $\theta \in \mathbb{R}^r$, which we call *flow functions*.

The module requires each of the following inputs:

1. a weighted, connected, undirected graph $G = (\mathcal{V}, \mathcal{E}, a)$ with edge weights $a \in \mathbb{R}^{|\mathcal{E}|}_{>0}$, and
2. a supply / demand vector $u \in \mathbb{R}^{|\mathcal{V}|}$ such that $\sum_{i \in \mathcal{V}} u_i = 0$.

The module outputs the unique vector $f \in \mathbb{R}^{|\mathcal{E}|}$ for which there exists $x \in \mathbb{R}^{|\mathcal{V}|}$ such that

$$Bf = u \tag{1}$$

$$f = [a]h_\theta(B^\top x) \tag{2}$$

where $B \in \{-1, 0, 1\}^{|\mathcal{V}| \times |\mathcal{E}|}$ is the signed incidence matrix of G, and h_θ is applied element-wise. We use the notation $\mathrm{FN}_{h,\theta}(G, u)$ to represent the solution f given inputs G and u, flow functions h, and parameters θ.

The IFN is illustrated in Figure 1. The IFN's only trainable component is its flow function, parameterized by θ. In practice, we will only make calls to the *inverse* of the flow function when evaluating

and backpropagating through IFN layers, so it is convenient to learn the inverse flow function directly. As the diagram highlights, IFNs are layers that can be situated in more complex architectures, with other models upstream estimating the supply / demand vector, edge weights, or even the topology. For example, in power systems, demand forecasting is a very well-studied problem [24, 25], and one can solve the economic dispatch problem to forecast power generation at each node [26], collectively leading to an estimate of the supply / demand vector.

Theorem 2.2 (Well-Posedness)**.** *IFNs are well-posed, i.e., there always exists an $f \in \mathbb{R}^m$ and $x \in \mathbb{R}^n$ that solve (1)–(2), and f is unique.*

We do not prove this theorem yet, as the existence and uniqueness is a trivial consequence of the theory to be developed later in this section.

2.1 Evaluating the Implicit Flow Network

Our approach to evaluating the implicit flow network is adapted from [27]. Any undirected graph G induces a direct decomposition of the space of flow vectors $\mathbb{R}^{|\mathcal{E}|}$: given the signed incidence matrix $B \in \{-1, 0, 1\}^{|\mathcal{V}|\times|\mathcal{E}|}$, the *cycle space* $\ker(B)$ and *cutset space* $\mathrm{Img}(B^\top)$ are orthogonal, and

$$\mathbb{R}^{|\mathcal{E}|} = \ker(B) \oplus \mathrm{Img}(B^\top)$$

We refer the reader to [28, §9.4] for a primer on cycle and cutset spaces. Accordingly, we decompose the vector $f = \mathrm{FN}_{h,\theta}(G, u)$ as $f = f_{\mathrm{cyc}} + f_{\mathrm{cut}}$, where $f_{\mathrm{cyc}} \in \ker(B)$ and $f_{\mathrm{cut}} \in \mathrm{Img}(B^\top)$. The cutset component is readily from (1), since $Bf = Bf_{\mathrm{cut}} = u$ implies that $f_{\mathrm{cut}} = B^\dagger u$. Then we must analyze (2) to solve for f_{cyc}. Define a *cycle projection matrix* $P \in \mathbb{R}^{m\times m}$ as the oblique projection onto $\ker(B)$ parallel to $\mathrm{Img}([a]B^\top)$:

$$P = I_m - [a]B^\top \left(B[a]B^\top\right)^\dagger B \tag{3}$$

Based on this projection, we define a map $T : \ker(B) \to \ker(B)$ for all $f_{\mathrm{cyc}} \in \ker(B)$ by

$$T(f_{\mathrm{cyc}}) = P\left(f_{\mathrm{cyc}} - d_{\min}[a]h_\theta^{-1}([a]^{-1}f_{\mathrm{cyc}} + [a]^{-1}B^\dagger u)\right) \tag{4}$$

One can show that f_{cyc} is the unique fixed point of T, and that T is a contraction mapping, leading to a simple algorithm to compute this fixed point.

Theorem 2.3 (Properties of T)**.** *Consider an implicit flow network with parameters $d_{\min}$, $d_{\max}$, and θ, with flow functions h_θ. Suppose that the inputs $G = (\mathcal{V}, \mathcal{E}, a)$ and $u \in \mathbb{1}_{|\mathcal{V}|}^\perp$ are given, and let $B \in \{-1, 0, 1\}^{|\mathcal{V}|\times|\mathcal{E}|}$ be the signed incidence matrix of G. The following are true:*

1. *T is a contraction mapping with respect to $||\cdot||_{2,[a]^{-\frac{1}{2}}}$, with Lipschitz constant*

$$\mathrm{Lip}(T) \le 1 - \frac{d_{\min}}{d_{\max}},$$

2. *the sequence of iterates $f_{\mathrm{cyc}}^{(k+1)} = T(f_{\mathrm{cyc}}^{(k)})$ starting from any initial condition $f_{\mathrm{cyc}}^{(0)} \in \ker(B)$ converges to a unique fixed point f_{cyc}, and*

3. *the output of the implicit flow network is given by*

$$\mathrm{FN}_{h,\theta}(G, u) = f_{\mathrm{cyc}} + B^\dagger u \tag{5}$$

Since f_{cyc} exists and is unique, Theorem 2.3 immediately proves the well-posedness of the implicit flow network (Theorem 2.2). Better yet, Theorem 2.3 provides a simple algorithm for computing this solution: pick any $f_{\mathrm{cyc}}^{(0)} \in \ker(B)$, repeatedly apply the map T until approximate convergence, then add $B^\dagger u$. Some care is required when implementing this map. Since P is a dense matrix with $|\mathcal{E}|^2$ entries, it is undesirable to explicitly construct the cycle space projection matrix for large networks. Instead, in order to project a vector $v \in \mathbb{R}^{|\mathcal{E}|}$, we can use the fact that

$$w \triangleq (B[a]B^\top)^\dagger Bv = \underset{w\in\mathbb{R}^n}{\mathrm{argmin}} \left\{||B[a]B^\top w - Bv||_2\right\}$$

so the projection is evaluated as $Pv = v - [a]B^\top w$. Using this method of projection to implement T, the fixed point iteration to compute $\mathrm{FN}_{h,\theta}(G, u)$ is stated in Algorithm 1.

Algorithm 1 Evaluating the implicit flow network.

1: $B \leftarrow$ signed incidence matrix of G
2: $f_{\text{cut}} \leftarrow \text{argmin}_{f_{\text{cut}} \in \mathbb{R}^m} \{||Bf_{\text{cut}} - u||_2\}$
3: $f_{\text{cyc}} \leftarrow \mathbb{0}_m$
4: $\Delta f_{\text{cyc}} \leftarrow \infty \mathbb{1}_m$
5: **while** $||\Delta f_{\text{cyc}}||_{2,[a]^{-\frac{1}{2}}} > \epsilon$ **do**
6: $\quad v \leftarrow d_{\min}[a]h_\theta^{-1}\left([a]^{-1}f_{\text{cyc}} + [a]^{-1}f_{\text{cut}}\right)$
7: $\quad w \leftarrow \text{argmin}_{w \in \mathbb{R}^n} \left\{||B[a]B^\top w - Bv||_2\right\}$
8: $\quad \Delta f_{\text{cyc}} \leftarrow v - [a]B^\top w$
9: $\quad f_{\text{cyc}} \leftarrow f_{\text{cyc}} - \Delta f_{\text{cyc}}$
10: **end while**
11: $f \leftarrow f_{\text{cyc}} + f_{\text{cut}}$
12: **return** f

Theorem 2.4 (Implicit Flow Networks, Forward Pass)**.** *Consider an implicit flow network with parameters $d_{\min}$, $d_{\max}$, and θ, with flow functions h_θ. Suppose that the inputs $G = (\mathcal{V}, \mathcal{E}, a)$ and $u \in \mathbb{1}_{|\mathcal{V}|}^\perp$ are given, and let $B \in \{-1, 0, 1\}^{|\mathcal{V}| \times |\mathcal{E}|}$ be the signed incidence matrix of G. The following are true of Algorithm 1, with a tolerance of $\epsilon > 0$:*

1. *for each iteration $k = 1, 2, \ldots$ of the loop, let $f_{\text{cyc}}^{(k)}$ represent the new value of f_{cyc} defined on line 9; and let $f_{\text{cyc}}^{(0)} = \mathbb{0}_m$. Then*
$$f_{\text{cyc}}^{(k+1)} = T(f_{\text{cyc}}^{(k)}), \quad \forall k \geq 0;$$
2. *the algorithm converges with at most k^* iterations of the while loop, where*
$$k^* = 1 + \left\lceil \frac{\log\left(d_{\min}^{-1}\rho^{-1}\epsilon\right)}{\log\left(1 - \frac{d_{\min}}{d_{\max}}\right)} \right\rceil \tag{6}$$
and $\rho = ||[a]h_\theta^{-1}([a]^{-1}B^\dagger u)||_{2,[a]^{-\frac{1}{2}}}$; and
3. *the algorithm returns $f \in \mathbb{R}^{|\mathcal{E}|}$, where*
$$||f - \text{FN}_{h,\theta}(G, u)||_{2,[a]^{-\frac{1}{2}}} \leq \left(\frac{d_{\max} - d_{\min}}{d_{\min}}\right)\epsilon \tag{7}$$

If evaluating h_θ^{-1} is sufficiently simple, then the most expensive step in the iteration is solving the ordinary least squares problem on line 7. Using a general-purpose solver, the complexity of this operation is roughly $O(|\mathcal{V}|^3)$. But $B[a]B^\top$ is a sparse Laplacian matrix, so we can use a specialized Laplacian solver that reduces the complexity to $O(|\mathcal{E}| \log^k |\mathcal{E}|)$ for some constant k [29].

2.2 Computing the Gradients

In order to train the flow function and any upstream models, it is necessary to backpropagate gradients through the IFN layer. We can perform this backward pass using implicit differentiation, and it turns out that the gradients of $\text{FN}_{h,\theta}(G, u)$ with respect to the parameters θ, a, and u can also be computed using Algorithm 1, i.e., by writing the gradient as the output of an auxiliary implicit flow network.

Theorem 2.5 (Gradients)**.** *Consider an implicit flow network with parameters $d_{\min}$, $d_{\max}$, and θ, with flow functions h_θ. Suppose that the inputs $G = (\mathcal{V}, \mathcal{E}, a)$ and $u \in \mathbb{1}_{|\mathcal{V}|}^\perp$ are given, and let $B \in \{-1, 0, 1\}^{|\mathcal{V}| \times |\mathcal{E}|}$ be the signed incidence matrix of G. Let $f = \text{FN}_{h,\theta}(G, u)$, and let w be a scalar entry of θ, a, or u. We can compute the derivatives $\frac{df}{dw}$ as follows.*

Define a vector of flow functions $g : \mathbb{R}^{|\mathcal{E}|} \to \mathbb{R}^{|\mathcal{E}|}$ by
$$g(\eta) = \mathcal{D}^{-1}\left(\eta - [a]^{-1}\frac{\partial v}{\partial w}\right), \quad \forall \eta \in \mathbb{R}^{|\mathcal{E}|} \tag{8}$$

Flow	Nodal Variable	$h(y) =$
DC Current	Voltage	y
DC Power	Voltage	y^2
AC Power (lossless)	Voltage Angle	$\sin(y)$
Water Flow Rate	Hydraulic Head	$\mathrm{sign}(y)\lvert y\rvert^{0.54}$
Mechanical Force Networks	Position	y

Table 1: Examples of physical flow networks, with the interpretation of the nodal variables, and the flow function in their constitutive relationship.

where $\mathcal{D} \in \mathbb{R}^{|\mathcal{E}|\times|\mathcal{E}|}$ is the diagonal matrix with entries

$$\mathcal{D}_{ee} = \left.\frac{dh_\theta^{-1}(y_e)}{dy_e}\right|_{y_e = a_e^{-1} f_e}, \quad \forall e \in \mathcal{E} \tag{9}$$

and $v = [a]h_\theta^{-1}([a]^{-1} f_{\mathrm{cyc}} + [a]^{-1} B^\dagger u)$. Then

$$\frac{df}{dw} = \mathrm{FN}_{g,\cdot}(G, \mathbb{0}_n) + B^\dagger \frac{du}{dw} \tag{10}$$

(We use the notation $\cdot$ in place of θ, since g has no trainable parameters.) Furthermore, the derivative constraint parameters $d_{\min}, d_{\max}$ from the original implicit flow network are valid for the new implicit flow network.

In other words, to compute the gradient with respect to a parameter, we perform a single evaluation of the implicit flow network.

In order to compute the derivatives with respect to some parameter or input w, we first evaluate the partial derivatives $\frac{\partial v}{\partial w}$ and the total derivatives $\frac{du}{dw}$. Then we construct the flow functions g according to (8), and solve an implicit flow network to find $\frac{df}{dw}$ according to (10). It is trivial to evaluate $\frac{du}{dw}$, but for convenience, we provide the values of $\frac{\partial v}{\partial w}$ below:

$$\begin{aligned}
\frac{\partial v}{\partial \theta_i} &= [a]\frac{dh_\theta^{-1}([a]^{-1} f)}{d\theta_i}, \quad \forall i \in \{1, 2, \ldots, r\} \\
\frac{\partial v}{\partial a_e} &= \mathrm{diag}\left(h_\theta^{-1}([a]^{-1} f) - [a]^{-1}\mathcal{D} f\right)_e, \quad \forall e \in \mathcal{E} \\
\frac{\partial v}{\partial u_i} &= \left(\mathcal{D} B^\dagger\right)^{(e)}, \quad \forall e \in \mathcal{E}
\end{aligned}$$

3 Comparison with Other Models

3.1 Physical Models

The equations (1)–(2) that specify IFN's output are inspired by the governing equations of several physical flow networks. The steady-states of many network systems are determined by a conservation law and a constitutive relationship. DC power networks are a familiar example: edge currents f are uniquely determined by Kirchoff's current law $Bf = u$ and Ohm's law $f = [r]^{-1} B^\top x$, where u are nodal current injections, r are resistances, and x are nodal voltages. In lossless AC networks, active power flows f satisfy the conservation of energy $Bf = u$ and the power flow equation $f = [a]\sin(B^\top x)$, where the edge weights a are functions of voltage magnitudes, line reactances, and tap ratios; and we re-interpret x as a vector of voltage angles [30, §6.4]. Equations of the form (1)–(2) show up in other models as well, including the Hazen-Williams formula in water distribution networks [31, Sec. 8.15], and discretizations of models in continuum mechanics [32, §16.9]. Table 1 lists several of these networks, the physical interpretation of the nodal variables x, and the form of the flow function h.

3.2 Bilevel Optimization

Both of the state-of-the-art methods for flow estimation, from [3] and [4], use an optimization problem to predict flows. After a suitable transformation to incorporate external flow injections u, we can

state this optimization problem as

$$\hat{f} = \underset{f \in \mathbb{R}^{|\mathcal{E}|}}{\operatorname{argmin}} \left\{ ||f||_{2,[q]}^2 + \lambda^2 ||Bf - u||_2^2 \text{ s.t. } f_e = \tilde{f}_e, \ \ \forall \text{ labeled edges } e \in \mathcal{E} \right\} \tag{11}$$

where $\lambda > 0$, and $q > \mathbb{0}_m$ is some vector of edge weights. In [3], $q = \mathbb{1}_m$, while [4] allows q to be the output of a neural network. IFN is not explicitly an optimization problem, but it can be cast as one that is similar (11):

Theorem 3.1 (Optimization Form of IFN)**.** *Consider an IFN with flow function h_θ. Suppose that the inputs $G = (\mathcal{V}, \mathcal{E}, a)$ and $u \in \mathbb{1}_{|\mathcal{V}|}^{\perp}$ are given, and let $B \in \{-1, 0, 1\}^{|\mathcal{V}| \times |\mathcal{E}|}$ be the signed incidence matrix of G. Then the IFN output can be stated as the solution of a convex optimization problem:*

$$\mathrm{FN}_{h,\theta}(G, u) = \underset{f \in \mathbb{R}^m}{\operatorname{argmin}} \left\{ \sum_{e \in \mathcal{E}} \int_0^{f_e} h_\theta^{-1}(a_e^{-1} z) \, dz \text{ s.t. } Bf = u \right\} \tag{12}$$

Theorem 3.1 can be interpreted as a nonlinear generalization of the Thomson principle from electrical circuits theory [33]. Interestingly, the theorem sets up a direct comparison between IFN and the models in [3] and [4]. If the flow function h_θ is the identity map, then (12) can be simplified as

$$\mathrm{FN}_{h,\theta}(G, u) = \underset{f \in \mathbb{R}^m}{\operatorname{argmin}} \left\{ \sum_{e \in \mathcal{E}} ||f||_{2,[a]^{-1}}^2 \text{ s.t. } Bf = u \right\} \tag{13}$$

Ignoring the constraints from labeled flows, we can interpret (11) as using a penalty method to approximate the output of an IFN with a linear flow function. Thus, we have three distinct differences between IFN and the optimization-based approaches:

1. IFN allows for a nonlinear flow function, while [3] and [4] implicitly assume a linear constitutive relationship;
2. IFN imposes flow conservation as a hard constraint rather than an approximate constraint (which is a limitation if u is uncertain); and
3. IFN does not incorporate flow measurements directly; rather, the model exploits these measurements during training to learn the proper flow function (and train any upstream models for the IFN inputs), making it less sensitive to noise in the labeled flows.

4 Models for Flow Functions

In order to implement an IFN, it is necessary to parameterize its inverse flow function h_θ^{-1}. Since the flow function is essentially a trainable activation function, i.e., a scalar nonlinearity, simple models are likely to be sufficient. The main difficulty with selecting a flow function is that its slope must be bounded by $d_{\min} \leq h'_\theta(y) \leq d_{\max}$ for all $y \in \mathbb{R}$. This section proposes a simple scalar nonlinearity that is guaranteed to respect arbitrary upper and lower bounds on its slope.

Definition 4.1 (Derivative-Constrained Perceptron)**.** Let $k \in \mathbb{Z}_{>0}$ be a hidden layer size, let $a, b, c \in \mathbb{R}^k$ be freely trainable parameters (encoded within the parameter vector θ), and let σ be a non-expansive activation. Let $p, q \geq 1$ such that $p^{-1} + q^{-1} = 1$, and let $\bar{d}_{\min} \leq \bar{d}_{\max} \in \mathbb{R}$. Then the *derivative-constrained perceptron* $\mathrm{N}(x, \theta)$ is the 3-layer neural network defined by

$$\bar{c}(\theta) = \left(1 - \frac{(||c||_p ||a||_q - 1)_+}{||c||_p ||a||_q}\right) c \tag{L1}$$

$$\mathrm{N}_0(x, \theta) = \bar{c}^\top(\theta) \sigma(ax + b) \tag{L2}$$

$$\mathrm{N}(x, \theta) = \left(\frac{\bar{d}_{\max} - \bar{d}_{\min}}{2}\right) \mathrm{N}_0(x, \theta) + \left(\frac{\bar{d}_{\max} + \bar{d}_{\min}}{2}\right) x \tag{L3}$$

Intuitively, (L1) re-scales c so that the perceptron in (L2) is guaranteed to be non-expansive in x. Then (L3) re-centers and re-scales the derivatives of the perceptron from the range $[-1, 1]$ to $[\bar{d}_{\min}, \bar{d}_{\max}]$.

Theorem 4.2 (Derivative-Constrained Perceptron)**.** *Let $\mathrm{N}(x, \theta)$ be a derivative-constrained perceptron with $\bar{d}_{\min} \leq \bar{d}_{\max} \in \mathbb{R}$. Then for all parameter values θ,*

$$\bar{d}_{\min} \leq \frac{d}{dx} \mathrm{N}(x, \theta) \leq \bar{d}_{\max}, \ \ \forall x \in \mathbb{R} \tag{14}$$

Note that the values $\bar{d}_{\min}, \bar{d}_{\max}$ in Definition 4.1 and Theorem 4.2 are distinct from the IFN parameters $d_{\min}, d_{\max}$. Since we parameterize the *inverse* flow function h_θ^{-1} in IFN, one shoudl set $\bar{d}_{\min} = d_{\max}^{-1}$ and $\bar{d}_{\max} = d_{\min}^{-1}$ to implement h_θ^{-1} with a derivative-constrained perceptron.

5 Numerical Experiments

5.1 Datasets

We selected 6 standard power network test cases. The first 4 test cases (IEEE-57, IEEE-118, IEEE-145, and IEEE-300) are synthetic transmission system test cases, while the remaining cases ACTIVSg200 and ACTIVSg500 are similar to the Illinois and South Carolina power grids, respectively [34]. Each test case contains the topology and electrical parameters of the power network, as well as baseline demands and power injections at each node. While branch resistances are typically small, we set them to zero to ensure lossless transmission lines. We used the MATPOWER toolbox [35] to solve the power flow equations, then recorded the active power flows on each branch (f), computed the net active power injections at each node (u), and selected relevant electrical parameters as edge attributes (series reactance, tap ratio, and voltage magnitude at the two incident nodes). See the supplementary material for details.

5.2 Models and Experiment Details

IFN Architecture In order to use the IFN layer to predict power flows, we created a two-layer model. The first layer estimates positive edge weights $a \in \mathbb{R}^{|\mathcal{E}|}$ according to

$$a_e = \exp\left(\mathrm{L}(z_e)\right), \qquad \forall e \in \mathcal{E}$$

where L is a linear module, and z_e is the log-transformed vector of edge attributes. The second layer is an IFN, with a derivative-constrained perceptron as the inverse flow function ($k = 128$, $d_{\min} = \frac{2}{5}$, $d_{\max} = 2$, and $p = q = \frac{1}{2}$), with a ReLU activation function. The output of the IFN is used as the predicted flow vector.

Baselines We compared the IFN model against four baselines. The minimum divergence method (*Div*) from [3] minimizes the nodal divergence $||Bf||_2^2$ and a regularization term $\lambda||f||_2^2$. The bilevel optimization methods from [4] replace the uniform regularizer with a weighted regularizer $||f||_{2,[q]}^2$, where q is a vector of weights. In *Bil-MLP* and *Bil-GCN*, q is the output of either a 2-layer MLP or GCN model with edge attributes as inputs (we use 64 nodes in each hidden layer with ReLU activations). Finally, in *Bil-True*, we specify q as the reciprocal of the coefficient in the linearized constitutive relationship for AC power networks, so that *Bil-True* approximates (13) with a as the ground-truth edge weight. See the supplementary material for details and hyperparameter settings.

All of the baselines assume that nodal divergence Bf should be approximately zero, but nodes in power networks inject and withdraw power according to the supply / demand vector u, resulting in nonzero divergence. Thus, when we evaluate the baselines, we transform the power network into a divergence-free network by introducing a "source node", adding an edge from the source node to all nodes in $\mathcal{V}$, and treating the entries of u as the flows along each corresponding virtual edge.

Training We trained all models in a Google Colab notebook, using the Adam optimizer. We used an initial learning rate of 0.01, which we found to have good performance for all of the models. Training was terminated when the training loss had not decreased for 10 epochs.

5.3 Predicting Missing Flows

We studied the transductive task of predicting unlabeled flows, given that some labeled flows in the same network are known. If the edges $\mathcal{E}$ are partitioned into a labeled set $\mathcal{E}_l$ and an unlabeled set $\mathcal{E}_u$, the task is to predict the missing flows $\{f_e : e \in \mathcal{E}_u\}$ given the labeled flows $\{f_e : e \in \mathcal{E}_l\}$. For each of the 6 networks, we randomly selected a fraction of the edges to be labeled edges, and we trained IFN on the labeled edges by minimizing the RMSE loss function

$$\ell_{\mathrm{rmse}} = \sqrt{\frac{1}{|\mathcal{E}_l|} \sum_{e \in \mathcal{E}} \left(f_e - \mathrm{FN}_{h,\theta}(G, u)_e\right)^2}$$

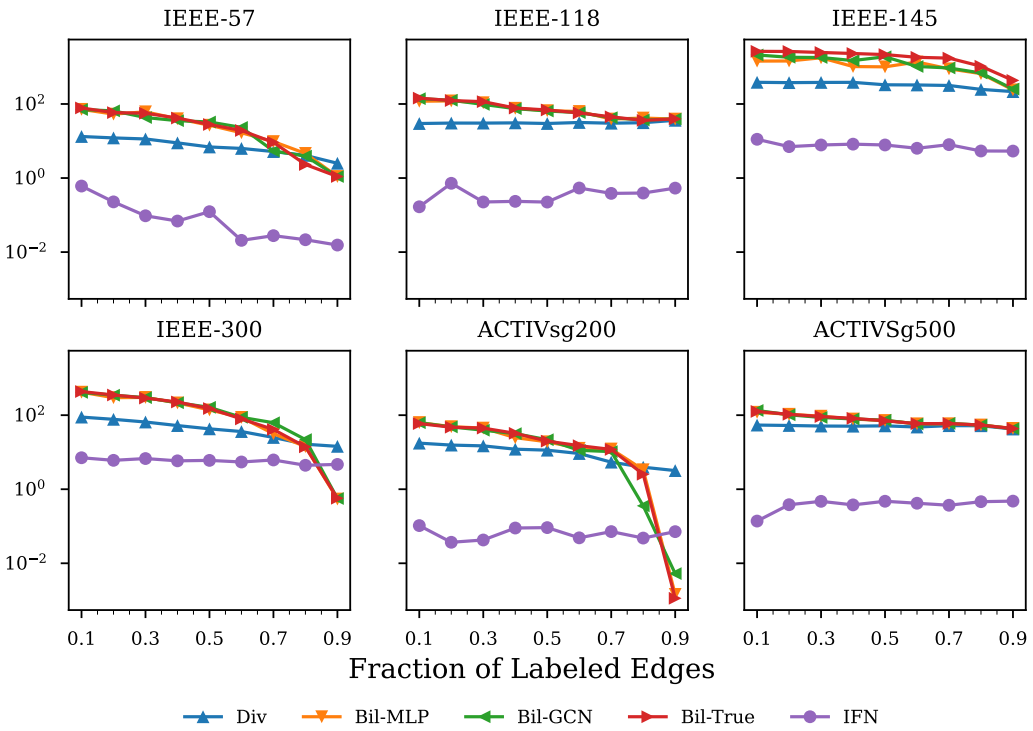

Figure 2: Results for missing flow prediction. Reported values are the RMSE (in units of MW) on the testing set, averaged across 10 trials. Note the vertical axis is in a logarithmic scale.

Then we evaluated the same function on the unlabeled edges $\mathcal{E}_u$ to compute the testing loss.

Figure 2 reports the results for various fractions of labeled edges. IFN significantly outperforms the baselines on all of the networks when less than 80% of the edges are labeled. While the other baselines tend to improve as more labeled edges are made available for training, IFN achieves near-optimal performance with as few as 10% of the edges labeled.

6 Conclusion

In this paper, we have introduced an implicit model for network flows that incorporates physics through a conservation law and through the existence of a constitutive relationship between flows and nodal variables. We have demonstrated that a simple architecture using this model can learn to accurately predict active power flows in AC networks. Future work may investigate more elaborate architectures using IFN as a layer, wherein the supply / demand vector, edge weights, or even the graph topology itself could be predicted from upstream models, and the flows themselves used for downstream tasks. Another interesting extension may be to apply our methodology to networks with higher-order interactions, i.e., hypergraphs [36].

IFN also has some limitations that should be addressed in future work. IFN assumes that the graph is undirected, which does not adequately model networks with unidirectional flows (e.g., traffic) or lossy flows (e.g., resistive power grids). IFN also assumes a constitutive relationship that depends on the *difference* between nodal variables. This form appears frequently in physical systems, but in other network flow models (like Daganzo traffic models [37]), the constitutive relationship has a more general dependence on the nodal variables. These limitations may be addressed with extensions of IFN's contraction algorithm.

References

[1] S. M. Rinaldi, J. P. Peerenboom, and T. K. Kelly. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Systems Magazine*, 21(6):11–25, 2001.

[2] J. A. Jacquez and C. P. Simon. Qualitative theory of compartmental systems. *SIAM Review*, 35(1):43–79, 1993.

[3] J. Jia, M. T. Schaub, S. Segarra, and A. R. Benson. Graph-based semi-supervised & active learning for edge flows. In *ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, page 761–771, July 2019.

[4] A. Silva, F. Kocayusufoglu, S. Jafarpour, F. Bullo, A. Swami, and A. K. Singh. Combining physics and machine learning for network flow estimation. In *International Conference on Learning Representations*, Online, May 2021.

[5] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. In *Canadian Journal of Mathematics*, volume 8, pages 399–404, 1956.

[6] M. Lippi, M. Bertini, and P. Frasconi. Short-term traffic flow forecasting: An experimental comparison of time-series analysis and supervised learning. *IEEE Transactions on Intelligent Transportation Systems*, 14(2):871–882, 2013.

[7] Y. Li, R. Yu, C. Shahabi, and Y. Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *International Conference on Learning Representations*, 2018.

[8] H. Yao, X. Tang, H. Wei, G. Zheng, and Z. Li. Revisiting spatial-temporal similarity: A deep learning framework for traffic prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 5668–5675, 2019.

[9] S. Bai, J. Z. Kolter, and V. Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems*, 2019.

[10] E. Winston and J. Z. Kolter. Monotone operator equilibrium networks. In *Advances in Neural Information Processing Systems*, 2020.

[11] M. Revay, R. Wang, and I. R. Manchester. Lipschitz bounded equilibrium networks. 2020.

[12] L. El Ghaoui, F. Gu, B. Travacca, A. Askari, and A. Tsai. Implicit deep learning. *SIAM Journal on Mathematics of Data Science*, 3(3):930–958, 2021.

[13] S. Jafarpour, A. Davydov, A. V. Proskurnikov, and F. Bullo. Robust implicit networks via non-Euclidean contractions. In *Advances in Neural Information Processing Systems*, December 2021.

[14] S. W. Fung, H. Heaton, Q. Li, D. McKenzie, S. Osher, and W. Yin. Fixed point networks: Implicit depth models with Jacobian-free backprop, 2021. ArXiv e-print.

[15] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, 2018.

[16] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter. Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems*, 2019.

[17] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

[18] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

[19] W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.

[20] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

[21] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pages 1263–1272. PMLR, 2017.

[22] P. Battaglia, R. Pascanu, M. Lai, D. Jimenez Rezende, and K. Kavukcuoglu. Interaction networks for learning about objects, relations and physics. In *Advances in Neural Information Processing Systems*. 2016.

[23] F. Gu, H. Chang, W. Zhu, S. Sojoudi, and L. El Ghaoui. Implicit graph neural networks. In *Advances in Neural Information Processing Systems*, 2020.

[24] A. K. Singh, Ibraheem, S. Khatoon, M. Muazzam, and D. K. Chaturvedi. Load forecasting techniques and methodologies: A review. In *2012 2nd International Conference on Power, Control and Embedded Systems*. IEEE, 2012.

[25] G. M. U. Din and A. K. Marnerides. Short term power load forecasting using deep neural networks. In *2017 International conference on computing, networking and communications (ICNC)*, pages 594–598. IEEE, 2017.

[26] K. D. Smith and K. Studarus. Limited-knowledge economic dispatch prediction using bayesian averaging of single-node models. In *2018 IEEE International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*. IEEE, 2018.

[27] S. Jafarpour, E. Y. Huang, K. D. Smith, and F. Bullo. Flow and elastic networks on the n-torus: Geometry, analysis and computation. *SIAM Review*, 64(1):59–104, 2022.

[28] F. Bullo. *Lectures on Network Systems*. Kindle Direct Publishing, 1.6 edition, January 2022.

[29] N. K. Vishnoi. $Lx = b$, Laplacian solvers and their algorithmic applications. *Theoretical Computer Science*, 8(1-2):1–141, 2013.

[30] P. Kundur. *Power System Stability and Control*. McGraw-Hill, 1994.

[31] E. John Finnemore and Joseph B. Franzini. *Fluid mechanics with engineering applications*. McGraw-Hill, tenth edition, 2002.

[32] J. R. Taylor. *Classical Mechanics*. University Science Books, 2005.

[33] P. G. Doyle and J. L. Snell. *Random Walks and Electric Networks*. Mathematical Association of America, 1984.

[34] A. B. Birchfield, T. Xu, K. M. Gegner, K. S. Shetye, and T. J. Overbye. Grid structural characteristics as validation criteria for synthetic networks. *IEEE Transactions on power systems*, 32(4):3258–3265, 2016.

[35] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas. MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Transactions on Power Systems*, 26(1):12–19, 2011.

[36] J. Jo, J. Baek, S. Lee, D. Kim, M. Kang, and S. J. Hwang. Edge representation learning with hypergraphs. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 7534–7546. Curran Associates, Inc., 2021.

[37] G. Como. On resilient control of dynamical flow networks. *Annual Reviews in Control*, 43:80–90, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] We do not forsee any negative societal impacts from flow estimation. On the contrary, we hope that these techniques will be useful for managing and protecting critical infrastructure.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Code and detailed instructions are included in the supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Training details are discussed in the supplementary material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] While error bars could be included in Figure 2, we felt it added considerable clutter to an already dense plot.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] As mentioned in Section 5.2, computation was performed in a Google Colab notebook. GPU type varies in Colab, so we did not report the specifics.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [No]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]