DELAY-AWARE BACKPRESSURE ROUTING USING GRAPH NEURAL NETWORKS

Zhongyuan Zhao⋆, Bojan Radojicic†, Gunjan Verma‡, Ananthram Swami‡, and Santiago Segarra⋆

⋆Rice University, USA   †University of Novi Sad, Serbia
‡US Army's DEVCOM Army Research Laboratory, USA

ABSTRACT

We propose a throughput-optimal biased backpressure (BP) algorithm for routing, where the bias is learned through a graph neural network that seeks to minimize end-to-end delay. Classical BP routing provides a simple yet powerful distributed solution for resource allocation in wireless multi-hop networks but has poor delay performance. A low-cost approach to improve this delay performance is to favor shorter paths by incorporating pre-defined biases in the BP computation, such as a bias based on the shortest path (hop) distance to the destination. In this work, we improve upon the widely-used metric of hop distance (and its variants) for the shortest path bias by introducing a bias based on the link duty cycle, which we predict using a graph convolutional neural network. Numerical results show that our approach can improve the delay performance compared to classical BP and existing BP alternatives based on *pre-defined* bias while being adaptive to interference density. In terms of complexity, our distributed implementation only introduces a one-time overhead (linear in the number of devices in the network) compared to classical BP, and a constant overhead compared to the lowest-complexity existing bias-based BP algorithms.

Index Terms— Backpressure routing, graph neural networks, scheduling duty cycle, independent set, bias, shortest path.

1. INTRODUCTION

Wireless multi-hop networks have been traditionally used in military communications, disaster relief, and wireless sensor networks, and are envisioned to support emerging applications such as connected vehicles, drone/robot swarms, xG (device-to-device, wireless backhaul, and non-terrestrial coverage), Internet of Things, and machine-to-machine communications [1–6]. An attractive feature of wireless multi-hop networks is its self-organizing capability without relying on infrastructure, enabled by distributed resource allocation schemes. Among those schemes, backpressure (BP) routing [7] is a well-established solution for resource allocation across the physical, media access control (MAC), and network layers [8–21]. In the BP algorithm, each node maintains a separate queue for packets to each destination (also denominated as commodity), routing decisions are made by selecting the commodity with the maximum differential
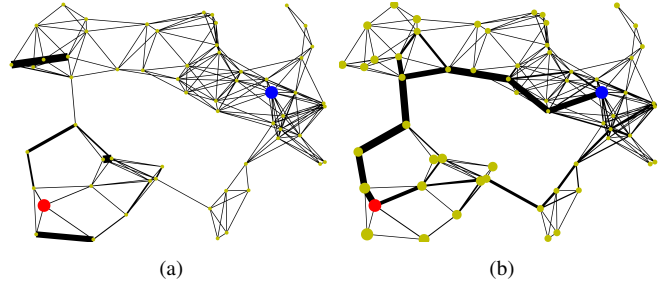
Fig. 1: A flow from the red node to the blue node in a wireless multi-hop network with 60 nodes. The normalized (across all links) number of packets sent over each link in 500 time slots is illustrated by their width. (a) Basic BP routing. (b) Enhanced dynamic BP routing (EDR) [8, 9] with a pre-defined bias (visualized by the node sizes) given by a scaled version of the shortest path to the destination.

queue backlog between the two ends of each link, and data transmissions are activated on a set of non-interfering links via MaxWeight scheduling [7]. The BP mechanism can drive the packets to explore all possible routes towards their destinations, while stabilizing the queues in the network for any flow rates within the network capacity region, i.e., BP achieves throughput optimality [7–9].

However, it is well-known that the classical BP routing suffers from poor delay performance for flows of low to medium rate [8–12], exhibiting undesirable characteristics such as *slow startup*, *random walk*, and the *last packet problem* [13, 14]. When a flow starts, many packets have to be first backlogged to form stable queue backlog-based gradients, causing large initial end-to-end delay. During BP scheduling, the fluctuations in queue backlogs drive packets towards random directions, causing unnecessarily long routes or loops. The phenomena of slow startup and random walk in BP routing are illustrated by the example in Fig. 1(a), in which packets from the red source node did not reach their blue destination in the first 500 time slots, but were trapped in two loops shown by the thickest edges.

The existing efforts to improve the delay performance of BP can be categorized into four types: 1) Use pre-defined queue-agnostic biases, such as the shortest path distance [8,9] or functions of it [10], to guide low-rate flows through the shortest routes. For example, with BP enhanced by the shortest path distance bias [8,9], in Fig. 1(b) packets can quickly reach their destination through two major routes around the empty central area. 2) Use queue-dependent biases that aggregate the queueing state information (QSI) of the local neighborhood (or global QSI) to improve the myopic BP decisions [11, 12] or use shadow queues [13, 15] to dynamically increase the backpressure. 3) Use delay metrics to replace the queue-dependent biases [11, 14, 16] or dynamically select routing schemes [18]. 4) Impose restrictions on the routes [17] or hop counts [19] to reduce the effect of loops. Solutions based on queue-agnostic biases are relatively simple, effective in mitigating the slow startup and random walk

problems, they are throughput optimal [8–10], and cost only a one-time communication overhead. In contrast, leveraging neighboring QSI incurs additional overhead per time slot, using delay metrics [11,14,16] and methods based on route restrictions [17,19] are more complex in implementation and could reduce the network capacity region. In addition, most of the aforementioned approaches require careful parameter tuning, typically done via trial-and-error. Besides those efforts, practical concerns in network connectivity [13,21] and network state uncertainty [20] have also been addressed.

In this work, our goal is to maintain the simplicity and computational tractability of the first type of existing works (queue-agnostic biases) but improve their performance through graph-based machine learning. Specifically, we seek to improve the widely-used bias given by the shortest hop distance to the destination [8–10,18,19], by considering the differences in link duty cycles in scheduling, rather than treating all links equally. We propose to predict the scheduling duty cycles of all the links with a graph neural network (GNN), an architecture that has been recently applied to improve the throughput [22,23], delay [24], and overhead [25] of link scheduling in wireless multi-hop networks. From the perspective of routing, the interfering wireless links are transformed into conflict-free links with effective rates equal to the link rates weighted by scheduling duty cycles, which can improve the routing decision, e.g., by avoiding hot-spots with many interfering neighbors (high betweenness centrality [26]). In addition, our approach can also lift the burden of parameter tuning by automatically increasing the weight of pre-defined biases in wireless networks with high interference degree.

Contribution. The contributions of this paper are twofold:
1) We propose an approach to leverage the scheduling duty cycle of links for delay-aware BP routing by learning appropriate biases using GNNs, and
2) Through numerical experiments, we demonstrate the superior delay performance of the shortest path bias based on link duty cycle in BP routing, especially in wireless networks with high interference.

2. SYSTEM MODEL

A wireless multi-hop network can be modeled as an undirected graph $\mathcal{G}^n = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is a set of nodes representing user devices in the network, and $\mathcal{E}$ represents a set of links, where $e = (i, j) \in \mathcal{E}$ for $i, j \in \mathcal{V}$ represents that node i and node j can talk to each other. We call $\mathcal{G}^n$ the connectivity graph, as it describes the connectivity relationship of the network. Here, we assume that $\mathcal{G}^n$ is a connected graph, so that two arbitrary nodes in the network can always reach each other. Notice that routing involves directed links, so we use $\overrightarrow{(i, j)}$ to denote data packets being transmitted from node i to node j over link (i, j). There is a set of flows $\mathcal{F}$ in the network, in which a flow $f = (i, c) \in \mathcal{F}$, where $i \neq c$ and $i, c \in \mathcal{V}$, describes the stream of packets from a source node i to a destination node c, potentially through multiple links. At a node $i \in \mathcal{V}$, there is a set of queues, $\{U_i^{(c)} | c \in \mathcal{V}\}$, in which $U_i^{(c)}$ represents the length of the queue for data packets destined to node c (or packets of commodity c).

To describe the conflict relationship between links, we define the *conflict graph*, $\mathcal{G}^c = (\mathcal{E}, \mathcal{C})$, as follows: a vertex $e \in \mathcal{E}$ represents a link in the network, and the presence of an undirected edge $(e_1, e_2) \in \mathcal{C}$ captures the interference relationship between links $e_1, e_2 \in \mathcal{E}$. We will be focusing on two popular models often used to define the conflict relationship between two links: 1) Interface conflict, where the conflict graph is given by the line graph of the connectivity graph. This represents the case where two links sharing the same node cannot be turned on simultaneously, e.g., if each node

is equipped with only one radio transceiver. 2) Physical distance interference model [27], which arises when two links interfere with each other if their incident users are within a certain distance such that their simultaneous transmission will cause the outage probability to exceed a prescribed level. A simplified scenario where all the users transmit at identical power levels with an omnidirectional antenna can be captured by the unit-disk interference model. In this model, two links conflict with each other if any of their nodes are within a pre-defined distance, which is the same for every pair of links. For the rest of this paper, we assume the conflict graph $\mathcal{G}^c$ to be known, e.g., by each link monitoring the wireless channel, or through more sophisticated estimation as in [28].

The MAC of the wireless network is assumed to be time-slotted orthogonal multiple access. Each time slot t contains a stage of decision making for routing and scheduling, followed by the second stage of data transmission. Therefore, we use $U_i^{(c)}(t)$ to describe the queue of commodity c at node i at the beginning of time slot t. The exogenous packet arrivals are collected by the non-negative integer matrix $\mathbf{A} \in \mathbb{Z}_+^{|\mathcal{F}| \times T}$, in which the element of row f and column t, $\mathbf{A}_{f,t}$, is the number of packets arriving at the source node of flow f at time slot t. Matrix $\mathbf{R} \in \mathbb{Z}_+^{|\mathcal{E}| \times T}$ collects the (stochastic) real-time link rates, of which an element $\mathbf{R}_{e,t}$ represents the number of packets that can be delivered over link e in time slot t. The long term link rate of a link $e \in \mathcal{E}$ is denoted by $r_e = \mathbb{E}_{t \leq T}[\mathbf{R}_{e,t}]$, and $\bar{r} = \mathbb{E}_{e \in \mathcal{E}, t \leq T}[\mathbf{R}_{e,t}]$ is the network-wide average link rate.

3. (BIASED) BACKPRESSURE ROUTING

Classical backpressure routing is a distributed algorithm for routing and scheduling, consisting of 4 steps. First, for each directed link $\overrightarrow{(i, j)}$, the BP algorithm selects the optimal commodity $c_{ij}^*(t)$ as the one with the maximal backpressure, which is defined as the difference of queue lengths between the sender i and the receiver j,

$$c_{ij}^*(t) = \underset{c \in \mathcal{V}}{\operatorname{argmax}}\{U_i^{(c)}(t) - U_j^{(c)}(t)\} . \tag{1}$$

In step 2, the maximum differential backlog of $\overrightarrow{(i, j)}$ is found as:

$$w_{ij}(t) = \max\{U_i^{(c_{ij}^*(t))}(t) - U_j^{(c_{ij}^*(t))}(t), 0\} . \tag{2}$$

In step 3, MaxWeight scheduling [7] finds the schedule $\mathbf{s}(t) \in \{0, 1\}^{|\mathcal{E}|}$ to activate a set of *non-conflicting links* achieving the maximum total utility, where the per-link utility is $u_{ij}(t) = \mathbf{R}_{ij,t}\tilde{w}_{ij}(t)$,

$$\mathbf{s}(t) = \underset{\tilde{\mathbf{s}}(t) \in \{0,1\}^{|\mathcal{E}|}}{\operatorname{argmax}} \tilde{\mathbf{s}}(t)^\top [\mathbf{R}_{*,t} \odot \tilde{\mathbf{w}}(t)] , \tag{3}$$

where vector $\mathbf{R}_{*,t}$ collects the real-time link rate of all links, vector $\tilde{\mathbf{w}}(t) = [\tilde{w}_{ij}(t) | (i, j) \in \mathcal{E}]$, where $\tilde{w}_{ij} = \max\{w_{ij}(t), w_{ji}(t)\}$, and the direction of the link selected by the max function will be recorded for step 4. Notice that the MaxWeight scheduling in (3) involves solving a maximum weighted independent set (MWIS) problem on the conflict graph (due to the implicit non-conflict constraint on the schedule), which is NP-hard [29]. Therefore, in practice, (3) is solved approximately by heuristics, such as centralized greedy maximal scheduler (GMS), distributed local greedy scheduler (LGS) [30], and GNN-enhanced schedulers [23]. In step 4, all of the real-time link rate $\mathbf{R}_{ij,t}$ of a scheduled link is allocated to its optimal commodity $c_{ij}^*(t)$. The final transmission and routing variables of commodity $c \in \mathcal{V}$ on link $\overrightarrow{(i, j)}$ is

$$\mu_{ij}^{(c)}(t) = \begin{cases} \mathbf{R}_{ij,t}, & \text{if } c = c_{ij}^*(t), w_{ij}(t) > 0, s_{ij}(t) = 1, \\ 0, & \text{otherwise.} \end{cases} \tag{4}$$

When a set of pre-defined biases for each pair of node and commodity, $\mathcal{B} = \{B_i^{(c)}|i, c \in \mathcal{V}\}$, are used to improve the delay performance [8–10], step 1 in (1) and step 2 in (2) now respectively become,

$$c_{ij}^*(t) = \operatorname*{argmax}_{c \in \mathcal{V}}\{\tilde{U}_i^{(c)}(t) - \tilde{U}_j^{(c)}(t)\} , \qquad (5)$$

$$w_{ij}(t) = \max\{\tilde{U}_i^{(c_{ij}^*(t))}(t) - \tilde{U}_j^{(c_{ij}^*(t))}(t), 0\} , \qquad (6)$$

where $\tilde{U}_i^{(c)}(t) = U_i^{(c)}(t) + B_i^{(c)}$. Since the bias $B_i^{(c)}$ is considered to be independent from the queue lengths, we can see it as a non-negative *constant*, i.e., it does not have to be updated in every time slot. In practice, $\mathcal{B}$ can still be updated from time to time to match the changing network topology.

4. LINK DUTY CYCLE PREDICTION WITH GNNS

In MaxWeight scheduling, the likelihood of a link being scheduled depends on its local conflict or interference topology and the network traffic load. We propose to use a GNN to predict the expected delay on each link, which can serve to design a better bias in (5)-(6) than simply the shortest path distance (in number of hops) between node i and node j for commodity c.

We propose to predict the link duty cycle $\mathbf{x} \in \mathbb{R}^{|\mathcal{E}|}$, as $\mathbf{x} = \Psi_{\mathcal{G}^c}(\mathbf{1}; \boldsymbol{\omega})$, where $\Psi_{\mathcal{G}^c}$ is an L-layered convolutional GNN defined on the conflict graph $\mathcal{G}^c$, and $\boldsymbol{\omega}$ is the collection of trainable parameters of the GNN. We define the output of an intermediate l-th layer of the convolutional GNN as $\mathbf{X}^l \in \mathbb{R}^{|\mathcal{E}| \times g_l}$, $\mathbf{X}^0 = \mathbf{1}^{|\mathcal{E}| \times 1}$, $\mathbf{x} = \mathbf{X}_{*0}^L$ (column 0 of $\mathbf{X}^L$), and the l-th layer of the GNN is expressed as

$$\mathbf{X}^l = \sigma_l\left(\mathbf{X}^{l-1}\boldsymbol{\Theta}_0^l + \boldsymbol{\mathcal{L}}\mathbf{X}^{l-1}\boldsymbol{\Theta}_1^l\right), \ l \in \{1, \ldots, L\}. \qquad (7)$$

In (7), $\boldsymbol{\mathcal{L}}$ is the normalized Laplacian of $\mathcal{G}^c$, $\boldsymbol{\Theta}_0^l, \boldsymbol{\Theta}_1^l \in \mathbb{R}^{g_{l-1} \times g_l}$ are trainable parameters (collected in $\boldsymbol{\omega}$), and $\sigma_l(\cdot)$ is the activation function of the l-th layer. The activation functions of the input and hidden layers are selected as leaky ReLUs, whereas a node-wise softmax activation is applied at the output layer (each row of $\mathbf{X}^L$). The input and output dimensions are set as $g_0 = 1$ and $g_L = 2$. Since $\boldsymbol{\mathcal{L}}$ in (7) is a local operator on $\mathcal{G}^c$, each row of $\mathbf{X}^l$, e.g., $\mathbf{X}_{e*}^l, e \in \mathcal{E}$ can be computed through neighborhood aggregation as the following local operation on link e,

$$\mathbf{X}_{e*}^l = \sigma_l\left(\mathbf{X}_{e*}^{l-1}\boldsymbol{\Theta}_0^l + \left[\mathbf{X}_{e*}^{l-1} - \sum_{u \in \mathcal{N}_{\mathcal{G}^c}(e)} \frac{\mathbf{X}_{u*}^{l-1}}{\sqrt{d(e)d(u)}}\right]\boldsymbol{\Theta}_1^l\right), \qquad (8)$$

where $\mathbf{X}_{e*}^l \in \mathbb{R}^{1 \times g_l}$ captures the lth-layer features on e, $\mathcal{N}_{\mathcal{G}^c}(e)$ denotes the set of (interfering) neighbors of e, and $d(\cdot)$ is the degree of a vertex in $\mathcal{G}^c$. Based on (8), the link duty cycle vector $\mathbf{x}$ can be computed in a fully distributed manner through L rounds of local message exchanges between $e \in \mathcal{E}$ and its neighbors, making our delay-enhanced BP routing a distributed algorithm.

With the estimated duty cycles of all the links, we propose two ways of setting the per-hop distance for finding the delay-aware shortest path between two nodes. Depending on availability of the long term link rate $r_e, e \in \mathcal{E}$, we define the per-hop delay distance of link e as $\delta_e = 1/x_e$ or $\delta_e = \bar{r}/(x_e r_e)$, where $\bar{r} = \mathbb{E}_{e \in \mathcal{E}, t \leq T}[\mathbf{R}_{e,t}]$ is the network-wide average link rate. By setting the edge weights of the connectivity graph $\mathcal{G}^n$ as $\boldsymbol{\delta} = [\delta_e|e \in \mathcal{E}]$, bias $B_i^{(c)}$ is set as the *weighted* shortest path distance between nodes i and c on $\mathcal{G}^n$. This distance can be computed via distributed algorithms for weighted single source shortest path (SSSP) when a node joins the network, or weighted all pairs shortest path (APSP) in general.

Complexity. For distributed implementation, the local communication complexity (defined as the rounds of local exchanges between a node and its neighbors) of the GNN is $\mathcal{O}(L)$. The distributed weighted SSSP with the Bellman-Ford algorithm [31, 32] and APSP with a very recent algorithm in [33] both take $\mathcal{O}(|\mathcal{V}|)$ rounds. Compared to the hop distance-based methods [8, 9], our approach bears additional $\mathcal{O}(L)$ rounds of communications (and larger message size). Notice that $\mathcal{B}$ can be reused over time slots until the network topology ($\mathcal{G}^n$ or $\mathcal{G}^c$) changes, which is critical for overhead reduction and scalability promotion.

Throughput Optimality. The classical BP algorithm can stabilize the queues in the network as long as the arrival rates of flows are within the network capacity region, which is proved by Lyapunov drift theory [8, 9]. The same proof also applies to BP with non-negative constant biases, as shown in [8, 9] and a similar proof in [10]. Therefore, our approach retains the throughput optimality of the classical BP, since our proposed delay-aware shortest path bias, $B_i^{(c)}$, is non-negative, based on link weight $\boldsymbol{\delta}_e > 0$, as $x_e > 0$ for all $e \in \mathcal{E}$ due to the softmax activation at the output layer in (7).

Training. The parameters $\boldsymbol{\omega}$ (collecting $\boldsymbol{\Theta}_0^l$ and $\boldsymbol{\Theta}_1^l$ across all layers l) of our GNN are trained on a set of routing instances defined on random network processes drawn from a target distribution Ω. More precisely, we draw several instances (indexed by k) of the network topology, flows, packet arrivals, and link rates ($\mathcal{G}^n(k)$, $\mathcal{G}^c(k), \mathcal{F}(k), \mathbf{A}(k), \mathbf{R}(k)) \sim \Omega$. For every instance, the GNN first predicts the link duty cycle vector $\mathbf{x}(k) = \Psi_{\mathcal{G}^c(k)}(\mathbf{1}; \boldsymbol{\omega})$, then biases $\mathcal{B}(k)$ are generated by the APSP algorithm based on the link distance vector $\boldsymbol{\delta}(k) = \operatorname{diag}^{-1}(\mathbf{x}(k))\mathbf{1}$, then we run the bias-based backpressure routing for T time slots, and collect the schedules for each time slot, $\mathbf{s}^k(t)$. We train the parameters $\boldsymbol{\omega}$ to minimize the following mean squared error loss

$$\ell(\boldsymbol{\omega}) = \mathbb{E}_\Omega\left(|\mathcal{E}|^{-1}\left\|\mathbf{X}^L(k) - \mathbb{E}_t\left(\left[\mathbf{s}^k(t); \mathbf{1} - \mathbf{s}^k(t)\right]\right)\right\|_2^2\right). \qquad (9)$$

Intuitively, by minimizing the loss in (9), we are choosing parameters $\boldsymbol{\omega}$ such that the softmax output of our GNN $\mathbf{X}^L$ is close to predicting the fraction of time that each link is scheduled. We do not seek to learn this complex function for *any specific* topology but rather we want to minimize the average error over instances drawn from Ω. In practice, we approximate the expected values in (9) with the corresponding empirical averages. Indeed, with the collected experience tuples, we update the parameters $\boldsymbol{\omega}$ of the GNN after each training instance, through batch training with random memory sampling, employing the Adam optimizer. Notice that, once trained, the GNN can be used to compute the delay-aware bias for previously unseen topologies without any retraining. As long as the new topology, arrival rates, and link rates are similar to those observed in Ω during training, we illustrate in the next section that the trained GNN generalizes to new instances observed during testing.

5. NUMERICAL EXPERIMENTS

We evaluate different biased BP routing algorithms on simulated wireless multi-hop networks. As explained in Section 4, we consider two variants of our proposed delay-aware bias where the weights of each link are respectively defined as $\delta_e = 1/x_e$ (denoted by SP-$1/x$) and $\delta_e = \bar{r}/(x_e r_e)$ (denoted by SP-$\bar{r}/(xr)$). Recall that x_e is the predicted link duty cycle and r_e the long-term link rate. The competing benchmarks include the basic BP, virtual queue-based BP (VBR) [10], EDR [8, 9] with a bias defined by the shortest hop distance, including $\delta = 1$ (SP-Hop) and $\delta = 10$ (EDR-10), and BP
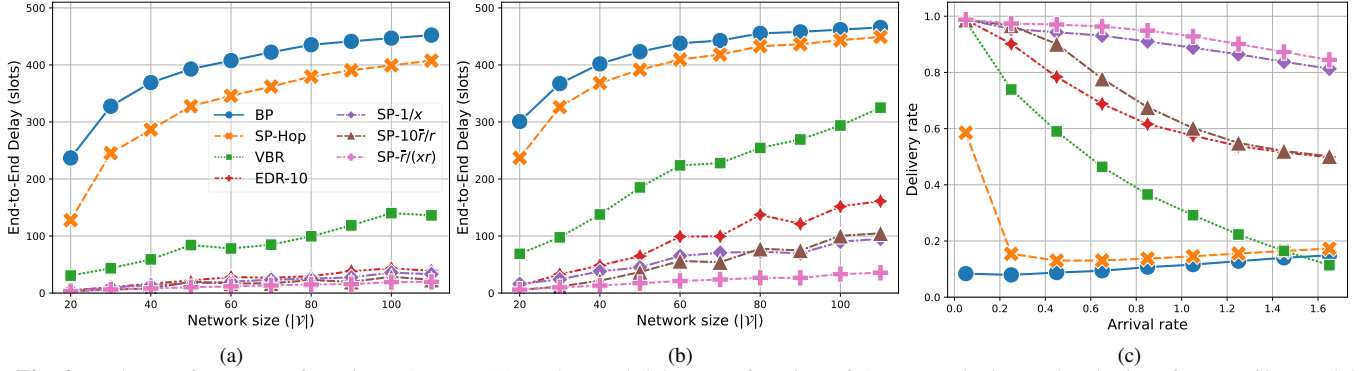
Fig. 2: Delay performance of routing schemes. (a) End-to-end delay as a function of the network size under the interface conflict model. (b) End-to-end delay as a function of the network size under unit-disk interference model. (c) Average delivery rate by $T = 1000$ as a function of the arrival rate on random networks of 100 nodes with unit-disk interference model.

based on shortest path bias defined by link rate, $\delta = 10\bar{r}/r$ (denoted by SP-$10\bar{r}/r$). The parameters of VBR and EDR are selected according to [10, 11].

The simulation is configured as follows. Every instance of a wireless multi-hop network is generated by a 2D point process with a given number of nodes $|\mathcal{V}| \in \{20, 30, \ldots, 110\}$ uniformly distributed in the plane with a constant density of $8/\pi$. A simplified scenario of a single channel is simulated. A link is established between two nodes if their distance is within 1.0, and the two conflict models previewed in Section 2 are considered: for the interface conflict model the conflict graph has an average degree of 12.4 whereas for the unit-disk interference model, the average conflict degree is 34.6. Notice that all the conflicts in the former model are included in the conflicts of the latter. Ten instances of such random networks are generated for each $|\mathcal{V}|$. On each randomly generated network, 10 random test instances are generated. Each test instance contains a number (uniformly chosen between $\lfloor 0.15|\mathcal{V}| \rfloor$ and $\lceil 0.30|\mathcal{V}| \rceil$) of random flows between different pairs of sources and destinations, where the exogenous packet arrival follows a Poisson process with a uniformly distributed arrival rate $\lambda(f) \sim \mathbb{U}(0.2, 1.0)$ for every $f \in \mathcal{F}$. Each test instance also includes a realization of uniformly distributed long-term link rates, $r_e \sim \mathbb{U}(10, 42)$, which are the expected values of the real-time link rates, $\mathbf{R}_{e,t} \sim \mathbb{N}(r_e, 3)$. The synthetic networks seek to represent wireless networks with uniformly distributed users of identical omnidirectional transmit power. The link rates are configured to capture fading channels with lognormal shadowing. Our proposed biases are generated based on a 5-layer convolutional GNN ($L = 5$, $g_l = 32, l \in \{1, \ldots, 4\}$), trained on a set of 100 random networks with $|\mathcal{V}| \in \{20, 30, \ldots, 60\}$, and the flows and link rates are configured similarly to the test settings.[1]

Our simulations lasts for $T = 1000$ time slots, where we compare the average end-to-end delay of packets attained by the 7 routing schemes considered. To be conservative, we treat the delay of an undelivered packet as $T - t_0$, where t_0 is the time it arrived at the source node. The estimated end-to-end delay of routing schemes as a function of the network size for both conflict models are presented in Figs. 2(a) and 2(b). The delays increase with network size for all methods due to the longer average hop distances of flows on larger topologies. The vanilla BP performs the worst due to its low delivery rate (fraction of packets delivered by the end of the simulation) that ranges from 0.56 in small networks ($|\mathcal{V}| = 20$) to 0.11 in larger networks ($|\mathcal{V}| = 110$). By introducing hop distance as bias, SP-Hop

can improve the delay of BP, and by scaling this bias by a factor of 10, EDR-10 can significantly reduce the delay of BP. The SP-$10\bar{r}/r$ method can further reduce delays of EDR-10 by $1/3 \sim 1/2$ through additional information of the long-term link rates across the network. Leveraging our GNN predictions of the scheduling duty cycles of links, SP-$1/x$ and SP-$\bar{r}/(xr)$ can further reduce the delays over EDR-10 and SP-$10\bar{r}/r$, respectively, where the improvements are more visible for the scenarios with higher average conflict degrees; see Fig 2(b). Contrary to the results in [10], the delay of VBR is much higher than EDR-10 in our tests, mainly due to the fact that VBR is designed for wireless sensor networks with only one sink, whereas flows in our simulation have different destinations. The test results show that our approach can generalize to networks larger than those seen during training, and can improve the delay of BP routing by offloading traffic from hot-spots congested by many interfering neighbors. Notice that shortest path biases (in hops) are myopic to this kind of information. Moreover, our GNN can learn the scale of the needed bias given the interference density of the network, alleviating the need of careful parameter tuning needed in the scaling of other methods, such as EDR-10 and SP-$10\bar{r}/r$.

Next, we evaluate the delay performance of routing schemes under different traffic loads, on 10 random networks of 100 nodes, under the unit-disk interference model. For each random network, 10 realizations of random flows and random link rates are generated like in the previous experiment, except that all the flows have identical arrival rate $\lambda \in \{0.05, 0.25, \ldots, 1.65\}$. The delivery rates of different routing schemes after $T = 1000$ time slots are presented in Fig. 2(c). Our approaches achieve the highest delivery rates under heavier network loads. This shows that we can improve both the delay and the delivery rate (number of packets delivered per unit of time) over existing biased BP algorithms across different traffic loads by leveraging our learning-based framework.

6. CONCLUSIONS

We improve the shortest path bias-based BP routing by replacing the well-established hop distance with a delay-aware distance. This delay-aware distance is computed using our predictions of the scheduling duty cycle of links given by a GNN. This solution inherits the simplicity, low per-slot overhead, and throughput optimality of BP routing with shortest path bias. Experiments show that our approach can outperform other BP algorithms based on pre-defined biases in terms of end-to-end delay and delivery rate across different interference densities and traffic loads, and exhibits good generalizability to larger networks.

---

[1]Training takes 5 hours on a workstation with a specification of 16GB memory, 8 cores, and Geforce GTX 1070 GPU. The source code is published at `https://github.com/zhongyuanzhao/dutyBP`

7. REFERENCES

[1] X. Lin, N. B. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1452–1463, 2006.

[2] S. K. Sarkar, T. G. Basavaraju, and C. Puttamadappa, *Ad hoc Mobile Wireless Networks: Principles, Protocols and Applications*. CRC Press, 2013.

[3] A. Kott, A. Swami, and B. J. West, "The internet of battle things," *Computer*, vol. 49, no. 12, pp. 70–75, 2016.

[4] I. F. Akyildiz, A. Kak, and S. Nie, "6G and beyond: The future of wireless communications systems," *IEEE Access*, vol. 8, pp. 133995–134030, 2020.

[5] "Cisco annual internet report (2018–2023)," white paper, Cisco Systems, Inc., Mar. 2020.

[6] X. Chen, D. W. K. Ng, W. Yu, E. G. Larsson, N. Al-Dhahir, and R. Schober, "Massive access for 5G and beyond," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 3, pp. 615–637, 2021.

[7] L. Tassiulas, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. on Automatic Control*, vol. 31, no. 12, 1992.

[8] M. Neely, E. Modiano, and C. Rohrs, "Dynamic power allocation and routing for time-varying wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, pp. 89–103, 2005.

[9] L. Georgiadis, M. J. Neely, L. Tassiulas, *et al.*, "Resource allocation and cross-layer control in wireless networks," *Foundations and Trends® in Networking*, vol. 1, no. 1, pp. 1–144, 2006.

[10] Z. Jiao, B. Zhang, W. Gong, and H. Mouftah, "A virtual queue-based back-pressure scheduling algorithm for wireless sensor networks," *EURASIP J. on Wireless Commun. and Netw.*, vol. 2015, no. 1, pp. 1–9, 2015.

[11] Y. Cui, E. M. Yeh, and R. Liu, "Enhancing the delay performance of dynamic backpressure algorithms," *IEEE/ACM Trans. Netw.*, vol. 24, no. 2, pp. 954–967, 2016.

[12] J. Gao, Y. Shen, M. Ito, and N. Shiratori, "Bias based general framework for delay reduction in backpressure routing algorithm," in *Intl. Conf. Computing, Networking and Communications (ICNC)*, pp. 215–219, IEEE, 2018.

[13] M. Alresaini, K.-L. Wright, B. Krishnamachari, and M. J. Neely, "Backpressure delay enhancement for encounter-based mobile networks while sustaining throughput optimality," *IEEE/ACM Trans. Netw.*, vol. 24, no. 2, pp. 1196–1208, 2016.

[14] B. Ji, C. Joo, and N. B. Shroff, "Delay-based back-pressure scheduling in multihop wireless networks," *IEEE/ACM Trans. Netw.*, vol. 21, no. 5, pp. 1539–1552, 2012.

[15] E. Athanasopoulou, L. X. Bui, T. Ji, R. Srikant, and A. Stolyar, "Back-pressure-based packet-by-packet adaptive routing in communication networks," *IEEE/ACM Trans. Netw.*, vol. 21, no. 1, pp. 244–257, 2012.

[16] L. Hai, Q. Gao, J. Wang, H. Zhuang, and P. Wang, "Delay-optimal back-pressure routing algorithm for multihop wireless networks," *IEEE Trans. Vehicular Tech.*, vol. 67, no. 3, pp. 2617–2630, 2017.

[17] A. Rai, C.-p. Li, G. Paschos, and E. Modiano, "Loop-free backpressure routing using link-reversal algorithms," *IEEE/ACM Trans. Netw.*, vol. 25, no. 5, pp. 2988–3002, 2017.

[18] P. Yin, S. Yang, J. Xu, J. Dai, and B. Lin, "Improving backpressure-based adaptive routing via incremental expansion of routing choices," in *ACM/IEEE Symp. Arch. for Networking and Communications Systems (ANCS)*, pp. 1–12, IEEE, 2017.

[19] L. Ying, S. Shakkottai, A. Reddy, and S. Liu, "On combining shortest-path and back-pressure routing over multihop wireless networks," *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 841–854, 2010.

[20] L. Ying and S. Shakkottai, "Scheduling in mobile ad hoc networks with topology and channel-state uncertainty," *IEEE Trans. on Automatic Control*, vol. 57, no. 10, pp. 2504–2517, 2012.

[21] J. Ryu, L. Ying, and S. Shakkottai, "Timescale decoupled routing and rate control in intermittently connected networks," *IEEE/ACM Trans. Netw.*, vol. 20, no. 4, pp. 1138–1151, 2012.

[22] Z. Zhao, G. Verma, C. Rao, A. Swami, and S. Segarra, "Distributed scheduling using graph neural networks," in *IEEE Int. Conf. on Acoustics, Speech and Signal Process. (ICASSP)*, pp. 4720–4724, 2021.

[23] Z. Zhao, G. Verma, C. Rao, A. Swami, and S. Segarra, "Link scheduling using graph neural networks," in *arXiv preprint arXiv:2109.05536*, 2022.

[24] Z. Zhao, G. Verma, A. Swami, and S. Segarra, "Delay-oriented distributed scheduling using graph neural networks," in *IEEE Int. Conf. on Acoustics, Speech and Signal Process. (ICASSP)*, pp. 8902–8906, 2022.

[25] Z. Zhao, A. Swami, and S. Segarra, "Distributed link sparsification for scalable scheduling using graph neural networks," in *IEEE Int. Conf. on Acoustics, Speech and Signal Process. (ICASSP)*, pp. 5308–5312, 2022.

[26] D. Katsaros, N. Dimokas, and L. Tassiulas, "Social network analysis concepts in the design of wireless ad hoc network protocols," *IEEE Network*, vol. 24, no. 6, pp. 23–29, 2010.

[27] W. Cheng, X. Cheng, T. Znati, X. Lu, and Z. Lu, "The complexity of channel scheduling in multi-radio multi-channel wireless networks," in *IEEE Intl. Conf. on Computer Comms. (INFOCOM)*, pp. 1512–1520, 2009.

[28] J. Yang, S. C. Draper, and R. Nowak, "Learning the interference graph of a wireless network," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 3, no. 3, pp. 631–646, 2016.

[29] C. Joo, G. Sharma, N. B. Shroff, and R. R. Mazumdar, "On the complexity of scheduling in wireless networks," *EURASIP J. on Wireless Commun. and Netw.*, vol. 2010, no. 1, p. 418934, 2010.

[30] C. Joo and N. B. Shroff, "Local greedy approximation for scheduling in multihop wireless networks," *IEEE Trans. on Mobile Computing*, vol. 11, no. 3, pp. 414–426, 2012.

[31] R. Bellman, "On a routing problem," *Quarterly of Applied Mathematics*, vol. 16, no. 1, pp. 87–90, 1958.

[32] L. R. Ford Jr, "Network flow theory," tech. rep., Rand Corp Santa Monica CA, 1956.

[33] A. Bernstein and D. Nanongkai, "Distributed exact weighted all-pairs shortest paths in near-linear time," in *ACM SIGACT Symp. Theory of Comp.*, pp. 334–342, 2019.