Throughput-Optimal Broadcast in Wireless Networks with Dynamic Topology

Abhishek Sinha, Leandros Tassiulas, *Fellow, IEEE*, and Eytan Modiano, *Fellow, IEEE*

Abstract—We consider the problem of throughput-optimal broadcasting in time-varying wireless network with an underlying Directed Acyclic (DAG) topology. Known broadcast algorithms route packets along pre-computed spanning trees. In large wireless networks with time-varying connectivities, the optimal trees are difficult to compute and maintain. In this paper we propose a new online throughput-optimal broadcast algorithm, which takes packet-by-packet scheduling and routing decisions, obviating the need for any global topological structures, such as spanning-trees. Our algorithm utilizes certain *queue-like* system-state information for making transmission decisions and hence, may be thought of as a generalization of the well-known *back-pressure algorithm*, which makes point-to-point unicast transmission decisions based on local queue-length information. Technically, the back-pressure algorithm is derived by greedily stabilizing the packet-queues. However, because of packet-duplications, the work-conservation principle is violated and appropriate queuing processes are difficult to define in the broadcast setting. To address this fundamental issue, we identify certain state-variables whose dynamics behave like *virtual queues*. By stochastically stabilizing these virtual queues, we devise a throughput-optimal broadcast policy. We also derive new characterizations of the broadcast-capacity of time-varying wireless DAGs and propose efficient algorithms to compute the capacity either exactly or approximately under various assumptions.

1 Introduction

The problem of efficiently disseminating packets from a source node to a subset of nodes in a network is known as the *Multicast problem*. In the special case when the incoming packets are to be distributed among all nodes in the network, the corresponding problem is referred to as the *Broadcast problem*. Multicasting and broadcasting are considered to be fundamental network functionalities, which enjoy numerous practical applications ranging from military communications [14], disaster management with mobile adhoc networks (MANET) [7], to streaming services for live web television [22].

There exists a substantial body of literature addressing different aspects of this problem in various networking settings. An extensive survey of various multicast routing protocols for MANET is provided in [10]. The authors of [6] consider the problem of minimum latency broadcast of a finite set of messages in MANET, where this problem is shown to be NP-hard. To address this issue, several approximation algorithms have been proposed in [9], all of which rely on construction of certain network-wide broadcast-trees. Cross-layer solutions for multi-hop multicasting in wireless network are given in [27] and [8]. These algorithms involve network coding, which introduces additional complexity and exacerbates end-to-end delay. The authors of [19] propose a multicast scheduling and routing protocol which balances load among a set of pre-computed spanning trees. However, these trees are difficult to compute and maintain in a large time-varying network. The authors of [24] propose a local control algorithm for broadcasting in a wireless network in the so called *scheduling-free model*, where an oracle is assumed to make interference-free scheduling decisions. This assumption, as noted by the authors themselves, is not practical.

In this paper we build upon our recent work in [21] and consider the problem of throughput-optimal broadcasting in a wireless network with time-varying connectivity. Throughout the paper, the overall network-topology is assumed to be a directed acyclic graph (DAG). We characterize the broadcast-capacity of time-varying wireless networks and propose an exact and an approximation algorithm to compute it efficiently. Next we propose a dynamic link-activation and packet-scheduling policy, which, unlike the previous algorithms, obviates the need to maintain any global topological structures, such as spanning trees, yet achieves the capacity. In addition to throughput-optimality, the proposed algorithm enjoys the attractive property of *in-order* packet-delivery, which makes it particularly useful in various online applications, e.g., VoIP and live multimedia communication [3]. Our algorithm is model-oblivious in the sense that its operation does not depend on detailed statistics of the random arrival or network-connectivity processes. We also show that the throughput-optimality of our algorithm is retained when the control decisions are made using *locally* available and possibly delayed state information.

Notwithstanding the vast literature on the general topic of broadcasting, to the best of our knowledge, this is the first work addressing throughput-optimal broadcasting in time-

- *A. Sinha and E. Modiano are with the Laboratory for Information and Decision Systems at MIT, Cambridge, MA. E-mail: {sinhaa, modiano}@mit.edu*
- *Leandros Tassiulas is with the dept. of Electrical Engg. and Yale Institute of Network Science, Yale University, New Haven, CT. E-mail: leandros.tassiulas@yale.edu*

A preliminary version of this paper appeared in the proceedings of ACM MobiHoc, 2016, held in Paderborn, Germany.
This work was supported by NSF Grant CNS-1217048, ONR Grant N00014-12-1-0064, and ARO MURI Grant W911NF-08-1-0238

varying wireless networks with store and forward routing. Our main technical contributions are as follows:

- We define and characterize the broadcast-capacity for wireless networks with time-varying connectivity. We show that the broadcast-capacity of time-varying wireless directed acyclic networks can be computed efficiently in some settings. We then derive tight upper and lower bounds on broadcast capacity and utilize it to propose an efficient approximation algorithm to estimate the capacity in a general setting.
- We propose a throughput-optimal dynamic routing and scheduling policy for broadcasting in a wireless DAG with time-varying connectivity. This algorithm is of *Max-Weight* type and critically uses the idea of *in-order* packet delivery. To the best of our knowledge, this is the first throughput-optimal dynamic algorithm proposed for broadcasting in time-varying wireless networks.
- We extend our algorithm to the practical scenario when the nodes have access only to delayed state information. We show that the throughput-optimality of the policy is retained even when the rate of inter-node communication is made arbitrarily small.
- We illustrate our theoretical findings with extensive numerical simulations.

The rest of the paper is organized as follows. Section 2 introduces the wireless network model. Section 3 defines and characterizes the broadcast capacity of a wireless DAG. It also provides an exact and an approximation algorithm to compute its broadcast-capacity. Section 4 describes our capacity-achieving broadcast algorithm. Section 5 extends the algorithm to the setting of imperfect state information. Section 6 provides numerical simulation results to illustrate our theoretical findings. In section 7 we summarize our results and conclude the paper.

2 Network Model

For pedagogical reasons, we first describe the model of a static wireless network without time-variation. Subsequently, we will incorporate time-variation into the static model. A static wireless network is modeled by a directed graph $\mathcal{G} = (V, E, \boldsymbol{c}, \mathcal{M})$, where V is the set of nodes and E is the set of *directed* point-to-point wireless links. There are a total of $|V| = n$ nodes and $|E| = m$ edges in the network. The vector $\boldsymbol{c} = (c_{ij})$ denotes packet transmission capacities of the links when they are activated and $\mathcal{M} \subset 2^{\{0,1\}^{|E|}}$ is the set of all edge-incidence vectors corresponding to the set of all feasible link-activations, complying with the given interference-constraints. The structure of the activation set $\mathcal{M}$ depends on the interference model, e.g., under the primary or node-exclusive interference model [12], $\mathcal{M}$ corresponds to the set of all *matchings* of the graph $\mathcal{G}$. Time is slotted and at time-slot t, any subset of links from the activation set $\mathcal{M}$ may be activated. Thus, at most c_{ij} packets can be transmitted in a slot from node i to node j, when link (i, j) is activated.

Let $\mathtt{r} \in V$ be the *source* node. At slot t, $A(t)$ packets arrive at the source, with mean $\mathbb{E}(A(t)) = \lambda$ and a finite second moment. The arrivals are assumed to be i.i.d. across slots. The broadcast problem is to efficiently disseminate the packets from the source to all nodes in the network.

2.1 Model of Time-varying Wireless Connectivity

Next we incorporate time-variation into our static model described above. In a wireless network, the channel-SINRs vary with time because of fading, shadowing and node mobility [25]. To take this random variation into account, we consider a simple ON-OFF channel model, where at each slot an individual link can be in either one of the two states, ON and OFF. In the ON state, a link (i, j), if activated, can transmit c_{ij} packets per slot, while in the OFF state it can not transmit any packet [1]. In other words, at any slot the entire network can be in any one configuration, out of finitely many possible configurations, denoted by the set Ξ. Each element $\sigma \in \Xi$ corresponds to a sub-graph $\mathcal{G}(V, E_\sigma) \subset \mathcal{G}(V, E)$, where $E_\sigma \subset E$ denotes the set of links that are ON at that slot. At every time-slot t, one of the configurations in the set Ξ is randomly realized. The network-configuration at time t is represented by the vector $\boldsymbol{\sigma}(t) \in \{0, 1\}^{|E|}$, where

$$\boldsymbol{\sigma}(e, t) = \begin{cases} 1, & \text{if } e \in E_{\boldsymbol{\sigma}(t)} \\ 0, & \text{otherwise.} \end{cases}$$

At the time-slot t, the network controller can only activate a set of non-interfering links from the set $E_{\boldsymbol{\sigma}(t)}$ that are ON. The network-configuration $\{\boldsymbol{\sigma}(t)\}_{t\geq 1}$ evolves according to a stationary ergodic process with the stationary distribution $\{p(\sigma)\}_{\sigma\in\Xi}$ [11], where

$$\sum_{\sigma\in\Xi} p(\sigma) = 1, \quad p(\sigma) > 0, \quad \forall \sigma \in \Xi \tag{1}$$

Since the underlying physical processes responsible for time-variation are often spatially-correlated [1], [17], the distribution of the link-states is assumed to possess an arbitrary joint-distribution. The detailed parameters of this process depend on the ambient physical environment, which is often difficult to measure. In particular, it is unrealistic to assume that the controller has knowledge of the statistical parameters of the process $\{\boldsymbol{\sigma}(t)\}_{t\geq 1}$. Fortunately, our proposed dynamic throughput-optimal broadcast policy does not require the statistical characterization of the configuration-process or its stationary-distribution $p(\boldsymbol{\sigma})$. This makes the policy robust and suitable for use in a dynamic setting.

Notations and Nomenclature:

In this section we briefly discuss the notations and conventions used throughout the paper. All vectors are assumed to be column vectors. For any set $\mathcal{X} \subset \mathbb{R}^k$, its convex-hull is denoted by $\text{conv}(\mathcal{X})$. Let $(U, V \setminus U)$ be a disjoint partition of the set of vertices of the graph, such that the source $\mathtt{r} \in U$ and $U \subsetneq V$. Such a partition will be called a *proper-partition*. To each proper partition corresponding to the node set U, associate the *proper-cut* vector $\boldsymbol{u} \in \mathbb{R}^m$, defined as follows:

$$\begin{aligned} \boldsymbol{u}_{i,j} &= c_{i,j} \quad \text{if } i \in U, j \in V \setminus U, \ (i,j) \in E \\ &= 0 \quad \text{otherwise} \end{aligned} \tag{2}$$

1. Generalization of the ON-OFF model, to multi-level discretization of link-capacity is straight-forward.

Denote the special, single-node proper-partitions by $U_j \equiv V \setminus \{j\}$, and the corresponding proper-cut vectors by $\boldsymbol{u}_j$, $\forall j \in V \setminus \{\mathtt{r}\}$. The set of all proper-cut vectors in the graph $\mathcal{G}$ is denoted by $\mathcal{U}$.

The *in-neighbour* set $\partial^{\text{in}}(j)$ of a node j is defined to be the set of all nodes $i \in V$ such that there is a directed edge $(i,j) \in E$. i.e.,

$$\partial^{\text{in}}(j) = \{i \in V : (i,j) \in E\} \tag{3}$$

Similarly, we define the *out-neighbour* set of a node j as

$$\partial^{\text{out}}(j) = \{i \in V : (j,i) \in E\} \tag{4}$$

For any two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ in $\mathbb{R}^m$, define the coordinate-wise product $\boldsymbol{z} \equiv \boldsymbol{x} \odot \boldsymbol{y}$ to be a vector in $\mathbb{R}^m$ such that $z_i = x_i y_i, 1 \leq i \leq m$.

For any set $\mathcal{S} \subset \mathbb{R}^m$ and any vector $\boldsymbol{v} \in \mathbb{R}^m$, the symbol $\boldsymbol{v} \odot \mathcal{S}$ denotes the set of all vectors obtained by the coordinate-wise product of the vector $\boldsymbol{v}$ and the elements of the set $\mathcal{S}$, i.e.,

$$\boldsymbol{v} \odot \mathcal{S} = \{\boldsymbol{y} \in \mathbb{R}^m : \boldsymbol{y} = \boldsymbol{v} \odot \boldsymbol{s}, \boldsymbol{s} \in \mathcal{S}\} \tag{5}$$

The usual dot product between two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^m$ is defined as: $\boldsymbol{x} \cdot \boldsymbol{y} = \sum_{i=1}^m x_i y_i$.

3 Definition and Characterization of Broadcast Capacity

Intuitively, a network supports a broadcast rate λ if there exists an admissible policy, under which all network nodes receive distinct packets at the rate λ. The broadcast-capacity of a network is defined as the maximally supportable broadcast rate. Formally, we consider a class Π of admissible policies, where each policy $\pi \in \Pi$ consists of a sequence of actions $\{\pi_t\}_{t \geq 1}$, executed at every slot t. Each action π_t consists of the following operations:

- The scheduler observes the current network-configuration $\sigma(t)$ and activates a subset of links by choosing a feasible activation vector $\boldsymbol{s}(t) \in \mathcal{M}_{\sigma(t)}$. Here $\mathcal{M}_\sigma$ denotes the set of all feasible link-activation vectors in the sub-graph $\mathcal{G}(V, E_\sigma)$, complying with the underlying interference constraints. Analytically, an element $\boldsymbol{s}$ from the set $\mathcal{M}_\sigma$ is represented by its m-dimensional binary incidence-vector. Thus, its e^{th} component $\boldsymbol{s}_e = 0$ if $e \notin E_\sigma$. In compact notation, $\mathcal{M}_\sigma = \boldsymbol{\sigma} \odot \mathcal{M}$.
- Each node i forwards a subset of packets (possibly empty) to its out-neighbour j, over an activated link $(i,j) \in \boldsymbol{\sigma}(t)$. The policy-class Π includes policies that may use all past and future information, and may forward any subset of packets over a link, subject to the per-slot link-capacity constraint.

To formally introduce the notion of broadcast capacity, we define the random variable $R_i^\pi(T)$ to be the number of distinct packets received by node i up to time T, under the action of a policy $\pi \in \Pi$. The time average $\liminf_{T\to\infty} R_i^\pi(T)/T$ is the rate of packet-reception at node i.

Definition 1. A policy $\pi \in \Pi$ is called a "broadcast policy of rate λ" if all nodes receive distinct packets at rate λ, i.e.,

$$\min_{i \in V} \liminf_{T\to\infty} \frac{1}{T} R_i^\pi(T) = \lambda, \quad \text{w.p. } 1 \tag{6}$$

where λ is the packet arrival rate at the source node $\mathtt{r}$.

Definition 2. The broadcast capacity λ^* of a network is defined to be the supremum of all arrival rates λ, for which there exists a broadcast policy $\pi \in \Pi$ of rate λ.

In the following subsection, we derive an upper-bound on broadcast-capacity λ^*, which immediately follows from the previous definition.

3.1 An Upper-bound on Broadcast Capacity

Consider a policy $\pi \in \Pi$ that achieves a broadcast rate of at least $\lambda^* - \epsilon$, for an $\epsilon > 0$. That such a policy exists, follows from the definition of the broadcast capacity λ^*.

Now consider any proper-cut U of the network $\mathcal{G}$. By the definition of a proper-cut, there exists a node $i \notin U$. Let $\boldsymbol{s}^\pi(t, \boldsymbol{\sigma}(t)) = (s_e^\pi(t, \boldsymbol{\sigma}(t)), e \in E)$ be the link-activation vector chosen by policy π in slot t, upon observing the current network-configuration $\boldsymbol{\sigma}(t)$. The maximum number of packets that can be transmitted across the cut U in slot t is upper-bounded by the total capacity of all activated links across the cut-set U, given by $\sum_{e \in E_U} c_e s_e^\pi(t, \boldsymbol{\sigma}(t))$. Hence, the number of distinct packets received by node i by time T is at most the total available capacity across the cut U up to time T, subject to link-activation decisions of the policy π. In other words,

$$R_i^\pi(T) \leq \sum_{t=1}^T \sum_{e \in E_U} c_e s_e^\pi(t, \boldsymbol{\sigma}(t)) = \boldsymbol{u} \cdot \sum_{t=1}^T \boldsymbol{s}^\pi(t, \boldsymbol{\sigma}(t)) \tag{7}$$

i.e.,

$$\frac{R_i^\pi(T)}{T} \leq \boldsymbol{u} \cdot \Big(\frac{1}{T} \sum_{t=1}^T \boldsymbol{s}^\pi(t, \boldsymbol{\sigma}(t))\Big),$$

where the cut-vector $\boldsymbol{u} \in \mathbb{R}^m$, corresponds to the cut-set U, as in Eqn.(2). It follows that,

$$\begin{aligned}\lambda^* - \epsilon \quad &\stackrel{(a)}{\leq} \quad \min_{j \in V} \liminf_{T\to\infty} \frac{R_j^\pi(T)}{T} \leq \liminf_{T\to\infty} \frac{R_i^\pi(T)}{T} \\ &\leq \quad \liminf_{T\to\infty} \boldsymbol{u} \cdot \Big(\frac{1}{T} \sum_{t=1}^T \boldsymbol{s}^\pi(t, \boldsymbol{\sigma}(t))\Big),\end{aligned} \tag{8}$$

where the inequality (a) follows from the fact that π is a broadcast policy of rate at least $\lambda^* - \epsilon$. Since the above inequality holds for all proper-cuts $\boldsymbol{u}$, we have

$$\lambda^* - \epsilon \leq \min_{\boldsymbol{u} \in \mathcal{U}} \liminf_{T\to\infty} \boldsymbol{u} \cdot \Big(\frac{1}{T} \sum_{t=1}^T \boldsymbol{s}^\pi(t, \boldsymbol{\sigma}(t))\Big) \tag{9}$$

The following technical lemma will prove to be useful for deriving an upper-bound on the broadcast-capacity.

> ***Lemma 1.*** For any policy $\pi \in \Pi$, and any proper-cut vector $\boldsymbol{u}$, there exist a collection of vectors $\left(\boldsymbol{\beta}_\sigma^\pi \in \text{conv}(\mathcal{M}_\sigma)\right)_{\sigma\in\Xi}$, such that, the following holds *a.s.*
>
> $$\min_{\boldsymbol{u}\in\mathcal{U}} \liminf_{T\to\infty} \boldsymbol{u}\cdot\left(\frac{1}{T}\sum_{t=1}^{T}\boldsymbol{s}^\pi(t,\boldsymbol{\sigma}(t))\right) = \min_{\boldsymbol{u}\in\mathcal{U}} \boldsymbol{u}\cdot\left(\sum_{\sigma\in\Xi} p(\sigma)\boldsymbol{\beta}_\sigma^\pi\right)$$

See Appendix 8.1 for the proof of this lemma. The above lemma essentially replaces the minimum cut-set bound of an arbitrary activations in (9), by the minimum cut-set bound of a stationary randomized activation. Combining Lemma 1 with Eqn. (9), we conclude that for any policy $\pi \in \Pi$ of rate at least $\lambda^* - \epsilon$, there exists a collection of vectors $\{\boldsymbol{\beta}_\sigma^\pi \in \text{conv}(\mathcal{M}_\sigma)\}_{\sigma\in\Xi}$ such that

$$\lambda^* - \epsilon \le \min_{\boldsymbol{u}\in\mathcal{U}} \boldsymbol{u}\cdot\left(\sum_{\sigma\in\Xi} p(\boldsymbol{\sigma})\boldsymbol{\beta}_\sigma^\pi\right) \tag{10}$$

Maximizing the RHS of Eqn. (10) over all vectors $\{\boldsymbol{\beta}_\sigma \in \text{conv}(\mathcal{M}_\sigma), \sigma\in\Xi\}$ and letting $\epsilon \searrow 0$, we have the following universal upper-bound on the broadcast capacity λ^*

$$\lambda^* \le \max_{\boldsymbol{\beta}_\sigma\in\text{conv}(\mathcal{M}_\sigma)} \min_{\boldsymbol{u}\in\mathcal{U}} \boldsymbol{u}\cdot\left(\sum_{\sigma\in\Xi} p(\boldsymbol{\sigma})\boldsymbol{\beta}_\sigma\right) \tag{11}$$

Specializing the above bound for single-node cuts of the form $\boldsymbol{U}_j = (V\setminus\{j\}) \to \{j\}, \forall j \in V\setminus\{\mathtt{r}\}$, we have the following upper-bound

$$\lambda^* \le \max_{\boldsymbol{\beta}_\sigma\in\text{conv}(\mathcal{M}_\sigma)} \min_{j\in V\setminus\{\mathtt{r}\}} \boldsymbol{u}_j\cdot\left(\sum_{\sigma\in\Xi} p(\boldsymbol{\sigma})\boldsymbol{\beta}_\sigma\right) \tag{12}$$

It will be shown in Section 4 that in a DAG, our throughput-optimal policy π^* achieves a broadcast-rate equal to the RHS of the bound (12). In particular, we have the following theorem

> ***Theorem 3.1.*** The broadcast-capacity λ^*_{DAG} of a time-varying wireless DAG is given by:
>
> $$\lambda^*_{\text{DAG}} = \max_{\boldsymbol{\beta}_\sigma\in conv(\mathcal{M}_\sigma),\sigma\in\Xi} \min_{j\in V\setminus\{\mathtt{r}\}} \boldsymbol{u}_j\cdot\left(\sum_{\sigma\in\Xi} p(\boldsymbol{\sigma})\boldsymbol{\beta}_\sigma\right) \tag{13}$$

The above theorem shows that for computing the broadcast-capacity of a wireless DAG, the minimum in the bound in (11) is attained by the single-node cuts $\boldsymbol{u}_j$.

3.2 An Illustrative Example of Capacity Computation

In this section, we work out a simple example to illustrate the previous results.

Consider the simple wireless network shown in Figure (1), with node $\mathtt{r}$ being the source. The possible network configurations $\sigma_i, i = 1,2,3,4$ are also shown. One packet can be transmitted over a link if it is ON. Moreover, since the links are assumed to be point-to-point, even if both the

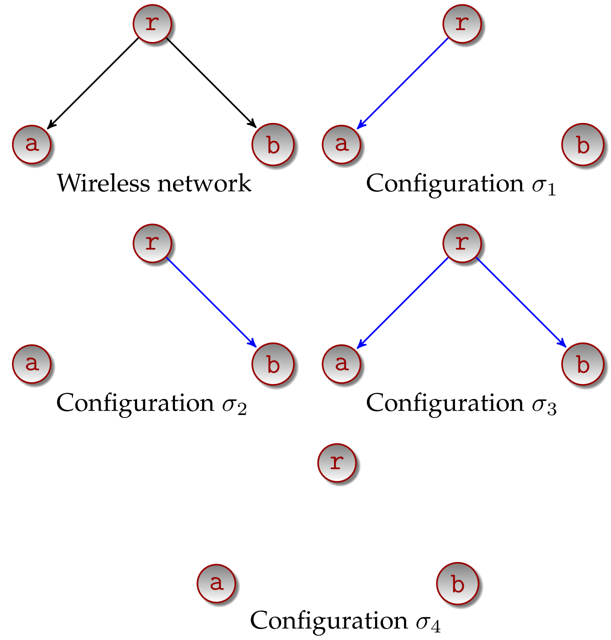

Fig. 1. A Wireless Network and its four possible configurations

links $\mathtt{ra}$ and $\mathtt{rb}$ are ON at a slot t (i.e., $\sigma(t) = \sigma_3$), a packet can be transmitted over one of the links only. Hence, the sets of feasible activations are given as follows:

$$\mathcal{M}_{\sigma_1} = \{\begin{pmatrix}1\\0\end{pmatrix}\}, \mathcal{M}_{\sigma_2} = \{\begin{pmatrix}0\\1\end{pmatrix}\},$$
$$\mathcal{M}_{\sigma_3} = \{\begin{pmatrix}1\\0\end{pmatrix}, \begin{pmatrix}0\\1\end{pmatrix}\}, \mathcal{M}_{\sigma_4} = \phi.$$

In the above vectors, the first coordinate corresponds to the edge $\mathtt{ra}$ and the second corresponds to the edge $\mathtt{rb}$.
To illustrate the effect of link-correlations on broadcast-capacity, we consider three different joint-distributions $p(\boldsymbol{\sigma})$, all of them having the identical marginal:

$$p(\mathtt{ra} = \text{ON}) = p(\mathtt{ra} = \text{OFF}) = \frac{1}{2}$$
$$p(\mathtt{rb} = \text{ON}) = p(\mathtt{rb} = \text{OFF}) = \frac{1}{2}$$

Case 1: Zero correlations: In this case, the links $\mathtt{ra}$ and $\mathtt{rb}$ are ON w.p. $\frac{1}{2}$ independently at every slot, i.e.,

$$p(\sigma_i) = 1/4, \quad i = 1,2,3,4 \tag{14}$$

It can be easily seen that the broadcast capacity, as given in Eqn. (13), is achieved when in configurations σ_1 and σ_2, the edges $\mathtt{ra}$ and $\mathtt{rb}$ are activated w.p. 1 respectively and in the configuration σ_3 the edges $\mathtt{ra}$ and $\mathtt{rb}$ are activated with probability $\frac{1}{2}$ and $\frac{1}{2}$. In other words, an optimal activation schedule of a corresponding stationary randomized policy is given as follows:

$$\boldsymbol{\beta}^*_{\sigma_1} = (1 \quad 0)', \boldsymbol{\beta}^*_{\sigma_2} = (0 \quad 1)', \boldsymbol{\beta}^*_{\sigma_3} = (\frac{1}{2} \quad \frac{1}{2})'$$

The optimal broadcast capacity can be computed from Eqn. (13) to be $\lambda^* = \frac{1}{4} + 0 + \frac{1}{4}\times\frac{1}{2} = \frac{3}{8}$.

Case 2: Positive correlations: In this case, the edges ra and rb are positively correlated, i.e., we have

$$p(\boldsymbol{\sigma}_1) = p(\boldsymbol{\sigma}_2) = 0;\ p(\boldsymbol{\sigma}_3) = p(\boldsymbol{\sigma}_4) = \frac{1}{2}$$

Then it is clear that half of the slots are wasted when both the links are OFF (i.e., in the configuration $\boldsymbol{\sigma}_4$). When the network is in configuration $\boldsymbol{\sigma}_3$, an optimal randomized activation is to choose one of the two links uniformly at random and send packets over it. Thus

$$\boldsymbol{\beta}^*_{\sigma_3} = (\frac{1}{2} \quad \frac{1}{2})'$$

The optimal broadcast-capacity, as computed from Eqn. (13) is $\lambda^* = \frac{1}{4}$.

Case 3: Negative correlations: In this case, the edges ra and rb are negatively correlated, i.e., we have

$$p(\boldsymbol{\sigma}_1) = p(\boldsymbol{\sigma}_2) = \frac{1}{2};\ p(\boldsymbol{\sigma}_3) = p(\boldsymbol{\sigma}_4) = 0$$

It is easy to see that, a capacity-achieving activation strategy in this case is to send packets over the link whichever is ON. The broadcast-capacity in this case is $\lambda^* = \frac{1}{2}$, the highest among the above three cases.

As apparent from the above example, with an arbitrary joint distribution of network-configurations $\{\boldsymbol{p}(\boldsymbol{\sigma})\}$, it is a matter of simple calculations to obtain the optimal activations $\boldsymbol{\beta}^*_{\boldsymbol{\sigma}}$ in Eqn. (13). However it is clear that for an arbitrary network with arbitrary activations $\mathcal{M}$ and configuration sets Ξ, evaluating (13) is non-trivial. The following section deals with this computational problem.

3.3 Efficient Computation of the Broadcast Capacity

In this section we study the problem of *efficient computation* of the Broadcast Capacity λ^* of a wireless DAG, given by Eqn. (13). In particular, we show that when the number of possible network configurations $|\Xi|(n)$ grows polynomially with n (the number of nodes in the network), there exists a strongly polynomial-time algorithm to compute λ^*, under the primary-interference constraint. Polynomially-bounded network-configurations arise, for example, when the set $\Xi(n)$ consists of subgraphs of the graph $\mathcal{G}$ with at most d number of edges, for some fixed integer d. In this case $|\Xi(n)|$ can be bounded as follows

$$|\Xi|(n) \leq \sum_{k=0}^{d} \binom{m}{k} = \mathcal{O}(n^{2d}),$$

where $m(= \mathcal{O}(n^2))$ is the number of edges in the graph $\mathcal{G}$.

Theorem 3.2 (Efficient Computation of λ^*). Suppose that for a wireless DAG $\mathcal{G}$ with n nodes, the number of possible network configurations $|\Xi|(n)$ is bounded polynomially in n. Then, there exists a strongly $\mathsf{polynomial-time}$ algorithm to compute the broadcast-capacity of the network under the primary interference constraints.

Although only polynomially many network configurations are allowed, we emphasize that Theorem (3.2) is highly non-trivial. This is because, each network-configuration $\sigma \in \Xi$ itself contains exponentially many possible activations (matchings) under the primary interference constraints. The key combinatorial result that leads to Theorem (3.2) is the existence of an efficient separator oracle for the matching-polytope for any arbitrary graph [20]. The detailed proof of Theorem 3.2 is provided below.

Proof: Under the primary interference constraint, the set of feasible activations of the graphs are *matchings* [26]. To solve for the optimal broadcast capacity in a time-varying network, we first rewrite the optimization problem involved in Eqn. (13) as a Linear Program (LP). Although this LP has exponentially many constraints, using a well-known separation oracle for matchings, we show that it is possible to solve this LP in strongly-polynomial time via the ellipsoid algorithm [2].

For a subset of edges $E' \subset E$, let $\chi^{E'}$ be the incidence vector, where $\chi^{E'}(e) = 1$ if $e \in E'$ and is zero otherwise. Let

$$\mathcal{P}_{\text{matching}}(\mathcal{G}(V,E)) = \textbf{convexhull}(\{\chi^M | M \text{ is a matching in } G(V,E)\})$$

We have the following classical result by Edmonds [20].

Theorem 3.3. The set $\mathcal{P}_{\text{matching}}(\mathcal{G}(V,E))$ is characterized by the set of all $\boldsymbol{\beta} \in \mathbb{R}^{|E|}$ such that :

$$\begin{aligned} \beta_e &\geq 0 \quad \forall e \in E \\ \sum_{e \in \partial^{\text{in}}(v) \cup \partial^{\text{out}}(v)} \beta_e &\leq 1 \quad \forall v \in V \\ \sum_{e \in E[U]} \beta_e &\leq \frac{|U|-1}{2}; \quad U \subset V,\ |U| \text{ odd} \end{aligned} \tag{15}$$

Here $E[U]$ is the set of edge with both end points in U.

Thus, following Eqn. (13), the broadcast capacity of a DAG can be obtained by the following LP :

$$\max \lambda \tag{16}$$

Subject to,

$$\lambda \leq \sum_{e \in \partial^{\text{in}}(v)} c_e \big(\sum_{\sigma \in \Xi} p(\sigma) \beta_{\sigma,e} \big),\ \forall v \in V \setminus \{r\} \tag{17}$$

$$\boldsymbol{\beta}_\sigma \in \mathcal{P}_{\text{matching}}(\mathcal{G}(V, E_\sigma)), \quad \forall \sigma \in \Xi \tag{18}$$

The constraint corresponding to $\sigma \in \Xi$ in (18) refers to the set of linear constraints given in Eqn.(15) corresponding to the graph $\mathcal{G}(V, E_\sigma)$, for each $\sigma \in \Xi$.

Invoking the equivalence of optimization and separation due to the ellipsoid algorithm [2], it follows that the LP (16) is solvable in poly-time, if there exists an efficient separator-oracle for the set of constraints (17) and (18). With our assumption of polynomially many network configurations $|\Xi|(n)$, there are only linearly many constraints ($n-1$, to be precise) in (17) with polynomially many variables in each constraint. Thus the set of constraints (17) can be separated efficiently. Next we invoke a classic result from the combinatorial-optimization literature which shows the existence of efficient separators for the matching polytopes.

> ***Theorem 3.4.*** [20] There exists a strongly poly-time algorithm, that given $\mathcal{G} = (V, E)$ and $\boldsymbol{\beta} : E \to \mathbb{R}^{|E|}$ determines if $\boldsymbol{\beta}$ satisfies (15) or outputs an inequality from (3.3) that is violated by $\boldsymbol{\beta}$.

Hence, there exists an efficient separator for each of the constraints in (3.3). Since there are only polynomially many network configurations, this directly leads to Theorem 3.2. □

3.4 Simple Bounds on λ^*

Using Theorem (3.2) we can, in principle, compute the broadcast-capacity λ^* of any wireless DAG with polynomially many network configurations. However, the complexity of the exact computation of λ^* grows substantially with the number of the possible configurations $|\Xi|(n)$. Moreover, Theorem (3.2) does not apply when $|\Xi|(n)$ can no longer be bounded by a polynomial in n. A simple example with exponentially large $|\Xi|(n)$ is the case when any link e is ON w.p. $p_e > 0$ i.i.d. at every slot.
To address this issue, we obtain bounds on λ^*, whose computational complexity is independent of the size of $|\Xi|$. These bounds are conveniently expressed in terms of the broadcast-capacity of the static network $\mathcal{G}(V, E)$ without time-variation, i.e. when $|\Xi| = 1$ and $E_\sigma = E, \sigma \in \Xi$. Let us denote the broadcast-capacity of the static network by λ^*_{stat}. Specializing Eqn. (13) to this case, we obtain

$$\lambda^*_{\text{stat}} = \max_{\boldsymbol{\beta} \in \text{conv}(\mathcal{M})} \min_{j \in V \setminus \{\mathtt{r}\}} \boldsymbol{u}_j \cdot \boldsymbol{\beta}. \tag{19}$$

Using Theorem (3.2), λ^*_{stat} can be computed in poly-time under the primary-interference constraint.
Now consider an arbitrary joint distribution $p(\boldsymbol{\sigma})$ such that each link is ON uniformly with probability p, i.e.,

$$\sum_{\boldsymbol{\sigma} \in \Xi : \boldsymbol{\sigma}(e) = 1} p(\boldsymbol{\sigma}) = p, \quad \forall e \in E. \tag{20}$$

We have the following bounds on λ^* for this case:

> ***Lemma 2 (Bounds on Broadcast Capacity).***
> $$p\lambda^*_{\text{stat}} \le \lambda^* \le \lambda^*_{\text{stat}}.$$

Proof: The proof consists of the following two parts.

3.4.1 Proof of the Upper-bound

Note that, for all $\sigma \in \Xi$, we have $E_\sigma \subset E$. Hence, it follows that

$$\mathcal{M}_\sigma \subset \mathcal{M}, \quad \forall \sigma \in \Xi$$

This in turn implies that

$$\beta_\sigma \in \text{conv}(\mathcal{M}_\sigma) \implies \beta_\sigma \in \text{conv}(\mathcal{M}) \tag{21}$$

Let an optimal solution to Eqn. (13) be obtained at $\big(\boldsymbol{\beta}^*_\sigma, \sigma \in \Xi\big)$. Then from Eqn. (21), it follows that

$$\sum_{\sigma \in \Xi} p(\boldsymbol{\sigma}) \boldsymbol{\beta}^*_\sigma \in \text{conv}(\mathcal{M})$$

Hence we have,

$$\max_{\boldsymbol{\beta}_\sigma \in \text{conv}(\mathcal{M}_\sigma)} \min_{j \in V \setminus \{\mathtt{r}\}} \boldsymbol{u}_j \cdot \big(\sum_{\sigma \in \Xi} p(\boldsymbol{\sigma}) \boldsymbol{\beta}^*_\sigma\big) \le \max_{\boldsymbol{\beta} \in \text{conv}(\mathcal{M})} \min_{j \in V \setminus \{\mathtt{r}\}} \boldsymbol{u}_j \cdot \boldsymbol{\beta}$$

Using Eqn. (19), this shows that

$$\lambda^* \le \lambda^*_{\text{stat}}$$

This proves the upper-bound.

3.4.2 Proof of the Lower-bound

Since $\mathcal{M}_\sigma \subset \mathcal{M}$, the expression for the broadcast-capacity (13) may be re-written as follows:

$$\lambda^* = \max_{\boldsymbol{\beta}_\sigma \in \mathcal{M}} \min_{j \in V \setminus \{\mathtt{r}\}} \sum_{e \in \partial^{\text{in}}(j)} c_e \big(\sum_{\sigma \in \Xi} p(\boldsymbol{\sigma}) \boldsymbol{\beta}_{\boldsymbol{\sigma}}(e) \mathbb{1}(e \in \boldsymbol{\sigma})\big)$$

Let $\boldsymbol{\beta}^* \in \text{conv}(\mathcal{M})$ be the optimal activation, achieving the RHS of (19). Hence we can lower-bound λ^* as follows

$$\begin{aligned} \lambda^* &\ge \min_{j \in V \setminus \{\mathtt{r}\}} \sum_{e \in \partial^{\text{in}}(j)} c_e \boldsymbol{\beta}^*(e) \big(\sum_{\sigma \in \Xi} p(\boldsymbol{\sigma}) \mathbb{1}(e \in \boldsymbol{\sigma})\big) \\ &\overset{(a)}{=} p \min_{j \in V \setminus \{\mathtt{r}\}} \sum_{e \in \partial^{\text{in}}(j)} c_e \boldsymbol{\beta}^*(e) \\ &= p \min_{j \in V \setminus \{\mathtt{r}\}} \boldsymbol{u}_j \cdot \boldsymbol{\beta}^* \\ &\overset{(b)}{=} p\lambda^*_{\text{stat}} \end{aligned}$$

Equality (a) follows from the assumption (20) and equality (b) follows from the characterization (19). This proves the lower-bound.

□

Generalization of the above Lemma to the setting, where the links are ON with non-uniform probabilities, may also be obtained in a similar fashion.
More importantly, as the example 3.2 shows, the simple bounds in Lemma 2 are tight. In this example the value of the connectivity parameter $p = \frac{1}{2}$, the lower-bound is attained in case (2) and the upper-bound is attained in case (3).
The above lemma immediately leads to the following corollary:

> ***Corollary 3.5.*** (APPROXIMATION-ALGORITHM FOR COMPUTING λ^*). Assume that, under the stationary distribution $p(\boldsymbol{\sigma})$, the probability that any link is ON is p, uniformly for all links. Then, there exists a poly-time p-approximation algorithm to compute the broadcast-capacity λ^* of a DAG, under the primary-interference constraints.

Proof: Consider the optimal randomized-activation vector $\boldsymbol{\beta}^* \in \text{conv}(\mathcal{M})$, corresponding to the stationary graph $\mathcal{G}(V, E)$ (19). By Theorem (3.2), $\boldsymbol{\beta}^*$ can be computed in poly-time under the primary interference constraint. Note that, by Caratheodory's theorem [15], the optimal $\boldsymbol{\beta}^*$ may be expressed as a convex combinbation of at most $|E|-1$ matchings. Thus it follows that λ^*_{stat} (19) may also be computed in poly-time.

From the proof of Lemma 2, it follows that by randomly activating $\boldsymbol{\beta}^*$ (i.e., $\boldsymbol{\beta}_{\boldsymbol{\sigma}}(e) = \boldsymbol{\beta}^*(e)\mathbb{1}(e \in \sigma), \forall \boldsymbol{\sigma} \in \Xi$) we obtain a broadcast-rate equal to $p\lambda^*_{\text{stat}}$ where λ^*_{stat} is shown to be an upper-bound to the broadcast capacity λ^* in Lemma (2). Hence it follows that $p\lambda^*_{\text{stat}}$ constitutes a p-approximation to the broadcast capacity λ^*, which can be computed in poly-time. □

This concludes our discussion the computational aspect of the broadcast-capacity. In the rest of the paper we are concerned with designing a dynamic and throughput-optimal broadcast policy for a time-varying wireless DAG network.

4 Throughput-Optimal Broadcast Policy for Wireless DAGs

The classical approach of solving the throughput-optimal broadcast problem in a static, wired network is to compute a set of edge-disjoint spanning trees of maximum cardinality (by invoking Edmonds' tree-packing theorem [18]) and then routing the incoming packets using these pre-computed trees [19]. In the time-varying wireless setting that we consider here, because of frequent and random changes in topology, routing packets over a fixed set of spanning trees is no-longer optimal. In particular, part of the network might become disconnected from time-to-time, and it is not clear how to select an optimal set of trees to disseminate packets. The problem becomes even more severe when the underlying statistical model of the network-connectivity process (in particular, the stationary distribution $\{p(\boldsymbol{\sigma}), \boldsymbol{\sigma} \in \Xi\}$) is unknown, which is often the case in mobile adhoc networks. Furthermore, wireless interference constraints add another layer of complexity, rendering the optimal dynamic broadcasting problem in wireless networks highly challenging.

In this section we propose an online, dynamic, throughput-optimal broadcast policy for time-varying wireless DAGs, that does not need to compute or maintain any global topological structures, such as spanning trees. Interestingly, we show that the broadcast-algorithm that we proposed earlier in [21] for static wireless networks, generalizes well to the time-varying case. As in [21], our algorithm also enjoys the attractive operational property *in-order* packet delivery. The key difference between the algorithm in [21] and our dynamic algorithm is in link-scheduling. In particular, in our algorithm, the activation sets are chosen based on current network-configuration $\boldsymbol{\sigma}(t)$.

4.1 Throughput-Optimal Broadcast Policy π^*

All policies $\pi \in \Pi$, that we consider in this paper, comprise of the following two sub-modules that are executed at every slot t:

- $\pi(\mathcal{A})$ (**Activation-module**): activates a subset of links $\boldsymbol{s}(t) \in \mathcal{M}_{\boldsymbol{\sigma}(t)}$, subject to the interference constraint and the current network-configuration $\boldsymbol{\sigma}(t)$.
- $\pi(\mathcal{S})$ (**Packet-Scheduling module**): schedules a subset of packets over the activated links.

Following our treatment in [21], we first restrict our attention to the policy sub-space $\Pi^{\text{in-order}}$, in which the admissible policies are required to follow the so-called *in-order* delivery property, defined as follows

Definition 3 (Policy-space $\Pi^{\text{in-order}}$ ***[21]).*** A policy π belongs to the space $\Pi^{\text{in-order}}$ if all incoming packets are serially indexed as $\{1, 2, 3, \ldots\}$ according to their order of arrival at the source $\mathtt{r}$ and a node can receive a packet p at time t, only if it has already received the packets $\{1, 2, , \ldots, p-1\}$.

As an immediate consequence of the *in-order* delivery property, the state of the received packets in the network at time-slot t may be succinctly represented by the n-dimensional vector $\boldsymbol{R}(t)$, whose i $^{\text{th}}$ component denotes the index of the *latest* packet received by node i by time t. We emphasize that this succinct network-state representation by the vector $\boldsymbol{R}(t)$ is valid only in the restricted policy-space $\Pi^{\text{in-order}}$. This compact state-representation results in substantial simplification of the overall state-space description. As a comparison, to completely specify the packet-configurations in the network at slot t in the general policy-space Π, we need to specify the *sets of packets* received by different nodes at slot t, which is quite unwieldy.

To effectively exploit the special structure of a DAG in designing our throughput-optimal broadcast policy, it will be useful to restrict our packet-scheduler $\pi(\mathcal{S})$ further to the following policy-space $\Pi^* \subset \Pi^{\text{in-order}}$.

Definition 4 (Policy-space $\Pi^* \subset \Pi^{\text{in-order}}$ ***[21]).*** A broadcast policy π belongs to the space Π^* if (1) $\pi \in \Pi^{\text{in-order}}$ and in addition (2) a packet p can be received by a node j at time t, only if all in-neighbours of the node j (i.e., nodes in $\partial^{\text{in}}(j)$) have received the packet p by the time t.

The above definition is illustrated in Figure 2. The variables $X_j(t)$ and $i^*_t(j)$ appearing in its description are defined subsequently in Eqn. (24).

It is easy to see that for all policies $\pi \in \Pi^*$, the packet

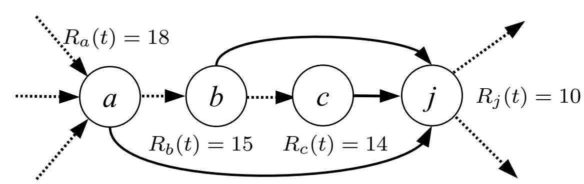

Fig. 2. Under a policy $\pi \in \Pi^*$, the set of packets available for transmission to node j at slot t is $\{11, 12, 13, 14\}$, which are available at all in-neighbors of node j. The in-neighbor of j inducing the smallest packet deficit is $i^*_t(j) = c$, and $X_j(t) = 14 - 10 = 4$.

scheduler $\pi(\mathcal{S})$ is *completely* specified. Hence, to specify a policy in the space Π^*, we need to define the activation-module $\pi(\mathcal{A})$ only.

Towards this end, let $\mu_{ij}(t)$ denote the rate (in packets per slot) allocated to the edge (i,j) in the slot t by a policy $\pi \in \Pi^*$. Note that, the allocated rate $\boldsymbol{\mu}(t)$ is constrained by the current network configuration $\boldsymbol{\sigma}(t)$ at slot t. In other words, we have

$$\boldsymbol{\mu}(t) \in \boldsymbol{c} \odot \mathcal{M}_{\boldsymbol{\sigma}(t)}. \tag{22}$$

This implies that, under any randomized activation

$$\mathbb{E}\boldsymbol{\mu}(t) \in \boldsymbol{c} \odot \mathrm{conv}(\mathcal{M}_{\boldsymbol{\sigma}(t)}). \tag{23}$$

In the following lemma, we show that for all policies $\pi \in \Pi^*$, certain state-variables $\boldsymbol{X}(t)$, derived from the state-vector $\boldsymbol{R}(t)$, satisfy the *Lindley recursions* [13] of queuing theory. Hence these variables may be thought of as *virtual queues*. This technical result will play a central role in deriving a *Max-Weight* type throughput-optimal policy π^*, which is obtained by stochastically stabilizing these virtual-queues. For each $j \in V \setminus \{\mathtt{r}\}$, define

$$X_j(t) = \min_{i \in \partial^{\text{in}}(j)} \big(R_i(t) - R_j(t)\big) \tag{24}$$

$$i_t^*(j) = \arg \min_{i \in \partial^{\text{in}}(j)} \big(R_i(t) - R_j(t)\big), \tag{25}$$

where in Eqn. (25), ties are broken lexicographically. The variable $X_j(t)$ denotes the minimum packet deficit of node j with respect to any of its in-neighbours. Hence, from the definition of the policy-space Π^*, it is clear that $X_j(t)$ is the maximum number of packets that a node j can receive from its in-neighbours at time t, under any policy in Π^*.
The following lemma proves a "queue-like-dynamics" of the variables $X_j(t)$, under any policy $\pi \in \Pi^*$.

Lemma 3 ([21]). Under any policy $\pi \in \Pi^*$, we have

$$X_j(t+1) \leq \Big(X_j(t) - \sum_{k \in \partial^{\text{in}}(j)} \mu_{kj}(t)\Big)^+ + \sum_{m \in \partial^{\text{in}}(i_t^*(j))} \mu_{m i_t^*(j)}(t) \tag{26}$$

Lemma (3) shows that the variables $\big(X_j(t), j \in V \setminus \{\mathtt{r}\}\big)$ satisfy Lindley recursions in the policy-space Π^*. Interestingly, unlike the corresponding unicast problem [23], there is no "physical queue" in the system.
Similar to the unicast problem [23], the next lemma shows that any activation module $\pi(\mathcal{A})$ that "stabilizes" the *virtual queues* $\boldsymbol{X}(t)$ for all arrival rates $\lambda < \lambda^*$, constitutes a throughput optimal broadcast-policy for a wireless DAG network.

Lemma 4. If under the action of a broadcast policy $\pi \in \Pi^*$, for all arrival rates $\lambda < \lambda^*$, the virtual queue process $\{\boldsymbol{X}(t)\}_0^\infty$ is rate-stable, i.e.,

$$\limsup_{T \to \infty} \frac{1}{T} \sum_{j \neq \mathtt{r}} X_j(T) = 0, \quad \text{w.p. } 1,$$

then the policy $\pi \in \Pi^*$ is a throughput-optimal broadcast policy for a wireless DAG network.

Proof: See Appendix 8.2. □

Equipped with Lemma (4), we now set out to derive a dynamic activation-module $\pi^*(\mathcal{A})$ to stabilize the virtual-queue process $\{\boldsymbol{X}(t)\}_0^\infty$ for all arrival rates $\lambda < \lambda^*$. Formally, the structure of the module $\pi^*(\mathcal{A})$ is defined by a mapping of the following form:

$$\pi^*(\mathcal{A}) : (\boldsymbol{X}(t), \boldsymbol{\sigma}(t)) \to \mathcal{M}_{\boldsymbol{\sigma}(t)}$$

Thus, the module $\pi^*(\mathcal{A})$ is stationary and dynamic as it depends on the current value of the state-variables and the network-configuration only. This activation-module is different from the policy described in [21] as the latter is meant for static wireless networks and hence, does not take into account the time-variation of network configurations, which is the focus of this paper.
To describe $\pi^*(\mathcal{A})$, we first define the following node-set

$$K_j(t) = \{m \in \partial^{\text{out}}(j) : j = i_t^*(m)\} \tag{27}$$

where the variables $i_t^*(m)$ are defined earlier in Eqn. (25). The activation-module $\pi^*(\mathcal{A})$ is given in Algorithm 1. The resulting policy in the space Π^* with the activation-module $\pi^*(\mathcal{A})$ is called π^*.

Algorithm 1 A Throughput-optimal Activation Module $\pi^*(\mathcal{A})$

1: To each link $(i,j) \in E$, assign a weight as follows:

$$W_{ij}(t) = \begin{cases} X_j(t) - \sum_{k \in K_j(t)} X_k(t), & \text{if } \boldsymbol{\sigma}_{(i,j)}(t) = 1 \\ 0, & \text{o.w.} \end{cases} \tag{28}$$

2: Select an activation $s^*(t) \in \mathcal{M}_{\boldsymbol{\sigma}(t)}$ as follows:

$$s^*(t) \in \arg \max_{\boldsymbol{s} \in \mathcal{M}_{\boldsymbol{\sigma}(t)}} \boldsymbol{s} \cdot \big(\boldsymbol{c} \odot \boldsymbol{W}(t)\big) \tag{29}$$

3: Allocate rates on the links as follows:

$$\boldsymbol{\mu}^*(t) = \boldsymbol{c} \odot \boldsymbol{s}^*(t) \tag{30}$$

Note that, in steps (1) and (2) above, the computation of link-weights and link-activations depend explicitly on the current network-configuration $\boldsymbol{\sigma}(t)$. As anticipated, in the following lemma, we show that the activation-module $\pi^*(\mathcal{A})$ stochastically stabilizes the virtual-queue process $\{\boldsymbol{X}(t)\}_0^\infty$.

Lemma 5. For all arrival rates $\lambda < \lambda^*$, under the action of the policy π^* in a DAG, the virtual-queue process $\{\boldsymbol{X}(t)\}_0^\infty$ is rate-stable, i.e.,

$$\limsup_{T \to \infty} \frac{1}{T} \sum_{j \neq \mathtt{r}} X_j(T) = 0, \quad \text{w.p. } 1$$

The proof of this lemma is based on a Lyapunov-drift argument [16]. Please refer to Appendix 8.3 for the complete proof.
Combining the lemmas (4) and (5), we immediately obtain the main result of this section

Theorem 4.1. The policy π^* is a throughput-optimal broadcast policy in a time-varying wireless DAG network.

5 Throughput-Optimal Broadcasting with Infrequent Inter-node Communication

In practical mobile wireless networks, it is unrealistic to assume that every node has perfect network-wide packet-state information at every slot. This is especially true in the case of time-varying dynamic networks, where network-connectivity changes frequently. In this section we extend the main results of section 4 by considering the setting where the nodes make control decisions with *imperfect* packet-state information that they currently possess. We will show that the dynamic broadcast-policy π^* retains its throughput-optimality even in this scenario.

5.0.0.1 State-Update Model: We assume that two nodes i and j can mutually update their knowledge of the set of packets received by the other node, only at those slots with positive probability, when the corresponding wireless-link (i,j) is in ON state. Otherwise, it continues working with the outdated packet state-information. Throughout this section, we assume that the nodes have perfect information about the current network-configuration $\boldsymbol{\sigma}(t)$.

Suppose that, the latest time prior to time t when packet-state update was made across the link (i,j) is $t-T_{(i,j)}(t)$. Here $T_{(i,j)}(t)$ is a random variable, supported on the set of non-negative integers. Assume that the network configuration process $\{\boldsymbol{\sigma}(t)\}_0^\infty$ evolves according to a finite-state, positive recurrent Markov-Chain, with the stationary distribution $\{p(\boldsymbol{\sigma})>0, \boldsymbol{\sigma}\in\Xi\}$. Using standard theory [5], it can be shown that the random variable $T(t)\equiv\sum_{(i,j)\in E}T_{(i,j)}(t)$ has bounded expectation for all time t .

5.0.0.2 Analysis of π^* with Imperfect Packet-State Information: Consider running the policy π^*, where each node j now computes the weights $W'_{ij}(t)$, given by Eqn.(28), of the in-coming links $(i,j)\in E$, based on the latest packet-state information available to it. In particular, for each of its in-neighbour $i\in\partial^{\text{in}}(j)$, the node j possess the following information of the number of packets received by node i:

$$R'_i(t)=R_i(t-T_{(ij)}(t)) \tag{31}$$

Now, if the packet-scheduler module $\pi'(\mathcal{S})$ of a broadcast-policy π' takes scheduling decision based on the imperfect state-information $\boldsymbol{R}'(t)$ (instead of the true state $\boldsymbol{R}(t)$), it still retains the following useful property:

Lemma 6. $\pi'\in\Pi^*$.

Proof: See Appendix 8.4. □

The above lemma states that the policy π' inherits the in-order delivery property and the in-neighbour packet delivery constraint of the policy-space Π^*.

From Eqn. (28) it follows that, computation of link-weights $\{W_{ij}(t), i\in\partial^{\text{in}}(j)\}$ by node j requires packet-state information of the nodes that are located within 2-hops from the node j. Thus, it is natural to expect that with an ergodic state-update process, the weights $W'_{ij}(t)$, computed from the imperfect packet-state information, will not differ too much from the true weights $W_{ij}(t)$, on the average. Indeed, we can bound the difference between the link-weights $W'_{ij}(t)$, used by policy π' and the true link-weights $W_{ij}(t)$, as follows:

Lemma 7. There exists a finite constant C such that, the expected weight $W'_{ij}(t)$ of the link (ij), locally computed by the node j using the random update process, differs from the true link-weight $W_{ij}(t)$ by at most C, i.e.

$$|\mathbb{E}W'_{ij}(t)-W_{ij}(t)|\leq C \tag{32}$$

The expectation above is taken with respect to the random packet-state update process.

Proof: See Appendix 8.5. □

From lemma (7) it follows that the policy π', in which link-weights are computed using imperfect packet-state information is also a throughput-optimal broadcast policy for a wireless DAG. Its proof is very similar to the proof of Theorem (4.1). However, since the policy π' makes scheduling decision using $\boldsymbol{W}'(t)$, instead of $\boldsymbol{W}(t)$, we need to appropriately bound the differences in drift using Lemma (7). The technical details are provided in Appendix 8.6.

Theorem 5.1. The policy π' is a throughput-optimal broadcast algorithm in a time-varying wireless DAG.

6 Numerical Simulation

We numerically simulate the performance of the proposed dynamic broadcast-policy on the 3×3 grid network, shown in Figure 3. All links are assumed to be of unit capacity. Wireless link activations are subject to primary interference constraints, i.e., at every slot, we may activate a subset of links which form a *Matching* [26] of the underlying topology. External packets arrive at the source node $\mathtt{r}$ according to a Poisson process of rate λ packets per slot. The following proposition shows that, the broadcast capacity λ^*_{stat} of the static 3×3 wireless grid (i.e., when all links are ON with probability 1 at every slot) is $\frac{2}{5}$.

Proposition 6.1. The broadcast-capacity λ^*_{stat} of the static 3×3 wireless grid-network in Figure 3 is $\frac{2}{5}$.

See Appendix 8.7 for the proof.

In our numerical simulation, the time-variation of the network is modeled as follows: link-states are assumed to

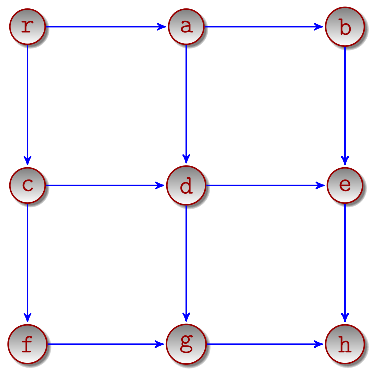

Fig. 3. A 3×3 grid network.

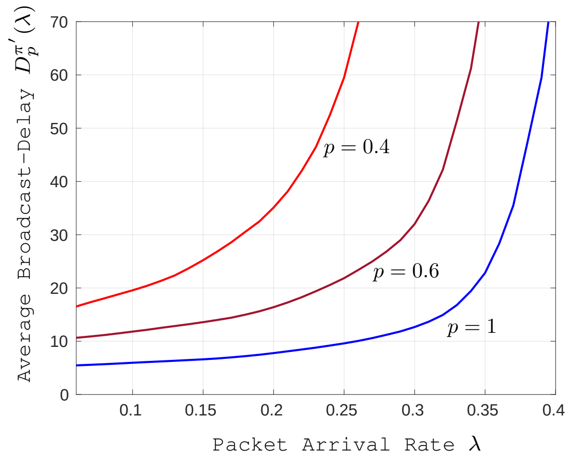

Fig. 4. Plot of average broadcast-delay $D_p^{\pi'}(\lambda)$, as a function of the packet arrival rates λ. The underlying wireless network is the 3×3 grid, shown in Figure 3, with primary interference constraints.

evolving in an i.i.d. fashion; each link is ON with probability p at every slot, independent of everything else. Here $0 < p \leq 1$ is the *connectivity-parameter* of the network. Thus, for $p = 1$ we recover the static network model of [21]. We also assume that the nodes have imperfect packet-state information as in Section 5. Two nodes i and j can directly exchange packet state-information, only when the link (i, j) (if any) is ON.

The average broadcast-delay $D_p^{\pi'}(\lambda)$ is plotted in Figure 4 as a function of the packet arrival rate λ. The broadcast-delay of a packet is defined as the number of slots the packet takes to reach all nodes in the network after its arrival. Because of the throughput-optimality of the policy π' (Theorem (5.1)), the broadcast-capacity $\lambda^*(p)$ of the network, for a given value of p, may be empirically evaluated from the λ-intercept of vertical asymptote of the $D_p^{\pi'}(\lambda) - \lambda$ curve.

As evident from the plot, for $p = 1$, the proposed dynamic algorithm achieves all broadcast rates below $\lambda^*_{\text{stat}} = \frac{2}{5} = 0.4$. This shows the throughput-optimality of the algorithm π'.

It is evident from the Figure 4 that the broadcast capacity $\lambda^*(p)$ is non-decreasing in the connectivity-parameter p, i.e., $\lambda^*(p_1) \geq \lambda^*(p_2)$ for $p_1 \geq p_2$. We observe that, with i.i.d. connectivity, the capacity bounds given in Lemma (2) are not tight, in general. Hence the lower-bound of $p\lambda^*_{\text{stat}}$ is a pessimistic estimate of the actual broadcast capacity $\lambda^*(p)$ of the DAG. The plot also reveals that, $D_p^{\pi'}(\lambda)$ is non-decreasing in λ for a fixed p and non-increasing in p for a fixed λ, as expected.

Variation of the Broadcast Capacity with increasing network size

To investigate the variation of the broadcast capacity of a dynamic wireless network with network size, we numerically simulate the proposed throughput-optimal broadcast policy π^* on three wireless grid networks with grid sizes 5×5, 8×8 and 9×9 respectively. The edge-connectivity processes are assumed to be i.i.d. with the connectivity parameter $p = 0.8$ uniformly over all edges. All simulated networks are assumed to be limited by primary interference constraints.

The time-averaged broadcast delay as a function of the incoming packet arrival rate is plotted in Figure 5 for the three simulated networks. The λ-intercepts of the vertical asymptotes of the delay-curves indicate the broadcast capacities $\lambda^*_{N \times N}$ of the corresponding grid networks. As expected, $\lambda^*_{N \times N}$ decreases with increasing network size.

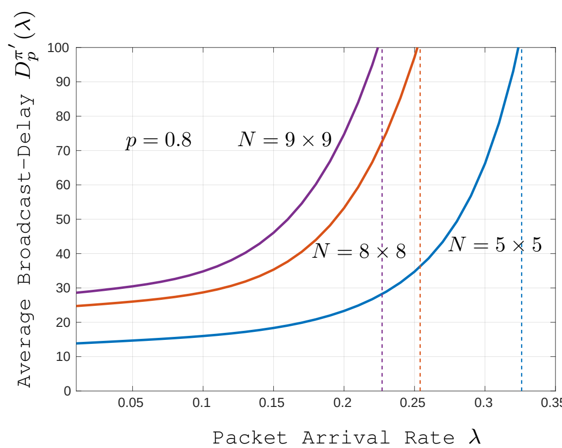

Fig. 5. Plot of average broadcast-delay $D_p^{\pi'}(\lambda)$, as a function of the packet arrival rates λ for i.i.d. connectivity process with parameter $p = 0.8$. Results are shown for 5×5, 8×8 and 9×9 grids with primary interference constraints. Vertical asymptotes indicate the respective broadcast-capacities of the networks.

7 Conclusion and Future Work

In this paper we studied the problem of throughput-optimal broadcasting in wireless directed acyclic networks with point-to-point links and time-varying connectivity. We characterized the broadcast-capacity of such networks and derived efficient algorithms for computing the same, both exactly and approximately. Next, we proposed a throughput-optimal broadcast policy for such networks. This policy does not need to maintain any spanning tree and operates based on locally available information, which is updated sporadically. The algorithm is robust and does not require statistics of the arrival or the connectivity process, thus making it useful for mobile wireless networks. The theoretical results are supplemented with illustrative numerical simulations. A possible future direction of research would be to remove the requirement of acyclic topology. It would also be interesting to extend the algorithm to wireless networks with point-to-multi-point links.

References

[1] P. Agrawal and N. Patwari. Correlated link shadow fading in multi-hop wireless networks. *Wireless Communications, IEEE Transactions on*, 8(8):4024–4036, 2009.

[2] D. Bertsimas and J. N. Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.

[3] Y. Chu, S. Rao, S. Seshan, and H. Zhang. Enabling conferencing applications on the internet using an overlay muilticast architecture. *ACM SIGCOMM computer communication review*, 31(4):55–67, 2001.

[4] S. Dasgupta, C. H. Papadimitriou, and U. Vazirani. *Algorithms*. McGraw-Hill, Inc., 2006.

[5] R. G. Gallager. *Discrete stochastic processes*, volume 321. Springer Science & Business Media, 2012.

[6] R. Gandhi, S. Parthasarathy, and A. Mishra. Minimizing broadcast latency and redundancy in ad hoc networks. In *Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, pages 222–232. ACM, 2003.

[7] M. Ge, S. V. Krishnamurthy, and M. Faloutsos. Overlay multicasting for ad hoc networks. In *Proceedings of Third Annual Mediterranean Ad Hoc Networking Workshop*, 2004.

[8] T. Ho and H. Viswanathan. Dynamic algorithms for multicast with intra-session network coding. In *Proc. 43rd Annual Allerton Conference on Communication, Control, and Computing*, 2005.

[9] S. C. Huang, P.-J. Wan, X. Jia, H. Du, and W. Shang. Minimum-latency broadcast scheduling in wireless ad hoc networks. In *26th IEEE INFOCOM 2007*.

[10] L. Junhai, Y. Danxia, X. Liu, and F. Mingyu. A survey of multicast routing protocols for mobile Ad-Hoc networks. *Communications Surveys Tutorials, IEEE*, 11(1):78–91, First 2009.

[11] A. Kamthe, M. A. Carreira-Perpiñán, and A. E. Cerpa. Improving wireless link simulation using multilevel markov models. *ACM Trans. Sen. Netw.*, 10(1):17:1–17:28, Dec. 2013.

[12] X. Lin, N. Shroff, and R. Srikant. A tutorial on cross-layer optimization in wireless networks. *Selected Areas in Communications, IEEE Journal on*, 24(8):1452–1463, Aug 2006.

[13] D. V. Lindley. The theory of queues with a single server. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 48, pages 277–289. Cambridge Univ Press, 1952.

[14] J. Macker, J. Klinker, and M. Corson. Reliable multicast data delivery for military networking. In *Military Communications Conference, 1996. MILCOM '96, Conference Proceedings, IEEE*, volume 2, pages 399–403 vol.2, Oct 1996.

[15] J. Matoušek. *Lectures on discrete geometry*, volume 108. Springer New York, 2002.

[16] M. J. Neely. Stochastic network optimization with application to communication and queueing systems. *Synthesis Lectures on Communication Networks*, 3(1):1–211, 2010.

[17] N. Patwari and P. Agrawal. Effects of correlated shadowing: Connectivity, localization, and rf tomography. In *Information Processing in Sensor Networks, 2008. IPSN'08. International Conference on*, pages 82–93. IEEE, 2008.

[18] R. Rustin. *Combinatorial Algorithms*. Algorithmics Press, 1973.

[19] S. Sarkar and L. Tassiulas. A framework for routing and congestion control for multicast information flows. *Information Theory, IEEE Transactions on*, 48(10):2690–2708, 2002.

[20] A. Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer, 2003.

[21] A. Sinha, G. Paschos, C. ping Li, and E. Modiano. Throughput-optimal broadcast on directed acyclic graphs. In *Computer Communications (INFOCOM), 2015 IEEE Conference on*, pages 1248–1256.

[22] D. Smith. IP TV Bandwidth Demand: Multicast and channel surfing. In *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 2546–2550, May 2007.

[23] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *Automatic Control, IEEE Transactions on*, 37(12):1936–1948, 1992.

[24] D. Towsley and A. Twigg. Rate-optimal decentralized broadcasting: the wireless case, 2008.

[25] D. Tse and P. Viswanath. *Fundamentals of wireless communication*. Cambridge university press, 2005.

[26] D. B. West et al. *Introduction to graph theory*, volume 2. Prentice hall Upper Saddle River, 2001.

[27] J. Yuan, Z. Li, W. Yu, and B. Li. A cross-layer optimization framework for multihop multicast in wireless mesh networks. *Selected Areas in Communications, IEEE Journal on*, 24(11), 2006.

8 Appendix

8.1 Proof of Lemma 1

Proof: Fix a time T. For each configuration $\sigma \in \Xi$, let $\{t_{\sigma,i}\}_{i=1}^{T_\sigma}$ be the index of the time-slots up to time T such that $\boldsymbol{\sigma}(t) = \boldsymbol{\sigma}$. Clearly we have,

$$\sum_{\sigma \in \Xi} T_\sigma = T \tag{33}$$

Hence, we can rewrite

$$\frac{1}{T}\sum_{t=1}^{T} \boldsymbol{s}^\pi(t, \boldsymbol{\sigma}(t)) = \sum_{\sigma\in\Xi} \frac{T_\sigma}{T}\frac{1}{T_\sigma}\sum_{i=1}^{T_\sigma} \boldsymbol{s}^\pi(t_{\sigma,i}, \boldsymbol{\sigma}) \tag{34}$$

Hence,

$$\boldsymbol{u}\cdot\Big(\frac{1}{T}\sum_{t=1}^{T} \boldsymbol{s}^\pi(t, \boldsymbol{\sigma}(t))\Big) = \sum_{\sigma\in\Xi} \frac{T_\sigma}{T}\boldsymbol{u}\cdot\Big(\frac{1}{T_\sigma}\sum_{i=1}^{T_\sigma} \boldsymbol{s}^\pi(t_{\sigma,i}, \boldsymbol{\sigma})\Big) \tag{35}$$

Since the process $\boldsymbol{\sigma}(t)$ is stationary ergodic, we have

$$\lim_{T\to\infty}\frac{T_\sigma}{T} = p(\sigma), \;\; \text{w.p. } 1 \;\; \forall \sigma \in \Xi \tag{36}$$

Using countability of Ξ and invoking the union bound, we can strengthen the above conclusion as follows

$$\lim_{T\to\infty}\frac{T_\sigma}{T} = p(\sigma), \;\; \forall \sigma \in \Xi, \;\; \text{w.p. } 1 \tag{37}$$

Hence from Eqn. (35) we have,

$$\begin{aligned}
&\min_{\boldsymbol{u}\in U}\liminf_{T\nearrow\infty} \boldsymbol{u}\cdot\Big(\frac{1}{T}\sum_{t=1}^{T}\boldsymbol{s}^\pi(t,\boldsymbol{\sigma}(t))\Big)\\
=\;&\min_{\boldsymbol{u}\in U}\sum_{\sigma\in\Xi} p(\boldsymbol{\sigma})\liminf_{T\to\infty}\boldsymbol{u}\cdot\Big(\frac{1}{T_\sigma}\sum_{i=1}^{T_\sigma}\boldsymbol{s}^\pi(t_{\sigma,i},\boldsymbol{\sigma})\Big), \quad \text{w.p. } 1
\end{aligned}$$

Since $p(\sigma) > 0, \forall \sigma \in \Xi$, the above implies that $T_\sigma \nearrow \infty$ as $T \nearrow \infty \forall \sigma$, w.p.1. In the rest of the proof we will concentrate on a typical sample path $\{\boldsymbol{\sigma}(t)\}_{t\geq 1}$ having the above property.
For each $\sigma \in \Xi$, define the sequence $\{\boldsymbol{\zeta}^\pi_{\sigma,T_\sigma}\}_{T_\sigma\geq 1}$

$$\boldsymbol{\zeta}^\pi_{\sigma,T_\sigma} = \frac{1}{T_\sigma}\sum_{i=1}^{T_\sigma}\boldsymbol{s}^\pi(t_{\sigma,i},\boldsymbol{\sigma}) \tag{38}$$

Since $\boldsymbol{s}^\pi(t_{\sigma,i},\boldsymbol{\sigma}) \in \mathcal{M}_\sigma$ for all $i \geq 1$, convexity of the set $\mathcal{M}_\sigma$ implies that $\boldsymbol{\zeta}^\pi_{\sigma,T_\sigma} \in \mathcal{M}_\sigma$ for all $T_\sigma \geq 1$. Since the set $\mathcal{M}_\sigma$ is closed and bounded (and hence, compact) any sequence in $\mathcal{M}_\sigma$ has a converging sub-sequence. Consider any set of converging sub-sequences $\{\boldsymbol{\zeta}^\pi_{\sigma,T_{\sigma_k}}\}_{k\geq 1}, \sigma \in \Xi$ such that, it achieves the following

$$\begin{aligned}
&\min_{u\in U}\sum_{\sigma\in\Xi} p(\sigma)\lim_{k\to\infty}\boldsymbol{u}\cdot\boldsymbol{\zeta}^\pi_{\sigma,T_{\sigma_k}}\\
=\;&\min_{u\in U}\sum_{\sigma\in\Xi} p(\sigma)\liminf_{T_\sigma\to\infty}\boldsymbol{u}\cdot\boldsymbol{\zeta}^\pi_{\sigma,T_\sigma}.
\end{aligned}$$

Let us denote

$$\lim_{k\to\infty} \boldsymbol{\zeta}^{\pi}_{\sigma,T_{\sigma_k}} = \boldsymbol{\beta}^{\pi}_{\sigma}, \quad \forall \sigma \in \Xi \tag{39}$$

Where $\boldsymbol{\beta}^{\pi}_{\sigma} \in \mathcal{M}_{\sigma}$, since $\mathcal{M}_{\sigma}$ is closed. Hence combining Eqn. (38), (39) and Eqn. (39), we have

$$\begin{aligned} & \min_{\boldsymbol{u}\in U} \liminf_{T\nearrow\infty} \boldsymbol{u}\cdot\left(\frac{1}{T}\sum_{t=1}^{T} \boldsymbol{s}^{\pi}(t,\boldsymbol{\sigma}(t))\right) \\ = \; & \min_{\boldsymbol{u}\in U} \sum_{\sigma\in\Xi} p(\boldsymbol{\sigma})\boldsymbol{u}\cdot\boldsymbol{\beta}^{\pi}_{\sigma} \\ = \; & \min_{\boldsymbol{u}\in U} \boldsymbol{u}\cdot\left(\sum_{\sigma\in\Xi} p(\boldsymbol{\sigma})\boldsymbol{\beta}^{\pi}_{\sigma}\right) \quad \text{w.p.1} \end{aligned}$$

□

8.2 Proof of Lemma (4)

Assume that under the policy $\pi \in \Pi^*$, the virtual queues $X_j(t)$ are rate stable i.e., $\lim_{T\to\infty} X_j(T)/T = 0, a.s.$ for all j. Applying union-bound, it follows that,

$$\lim_{T\to\infty} \frac{\sum_{j\neq r} X_j(T)}{T} = 0, \qquad \text{w.p. } 1 \tag{40}$$

Now consider any node $j \neq \mathtt{r}$ in the network. We can construct a simple path $p(\mathtt{r} = u_k \to u_{k-1} \ldots \to u_1 = j)$ from the source node $\mathtt{r}$ to the node j by running the following `Path construction algorithm` on the underlying graph $\mathcal{G}(V,E)$.

Algorithm 2 $\mathtt{r} \to j$ `Path Construction Algorithm`

Require: DAG $\mathcal{G}(V,E)$, node $j \in V$
1: $i \leftarrow 1$
2: $u_i \leftarrow j$
3: **while** $u_i \neq r$ **do**
4: $\quad u_{i+1} \leftarrow i^*_t(u_i)$;
5: $\quad i \leftarrow i+1$
6: **end while**

At time t, the algorithm chooses the parent of a node u_i in the path p as the one that has the least relative packet deficit as compared to u_i (i.e. $u_{i+1} = i^*_t(u_i)$). Since the underlying graph $\mathcal{G}(V,E)$ is a connected DAG (i.e., there is a path from the source to every other node in the network), the above path construction algorithm always terminates with a path $p(\mathtt{r} \to j)$. Note that the output path of the algorithm varies with time.

The number of distinct packets received by node j up to time T can be written as a telescoping sum of relative packet deficits along the path p, i.e.,

$$\begin{aligned} R_j(T) &= R_{u_1}(T) \\ &= \sum_{i=1}^{k-1} \left(R_{u_i}(T) - R_{u_{i+1}}(T)\right) + R_{u_k}(T) \\ &= -\sum_{i=1}^{k-1} X_{u_i}(T) + R_{\mathtt{r}}(T) \\ &\overset{(a)}{=} -\sum_{i=1}^{k-1} X_{u_i}(T) + \sum_{t=0}^{T-1} A(t), \end{aligned} \tag{41}$$

where the equality (a) follows the observation that

$$X_{u_i}(T) = Q_{u_{i+1}u_i}(T) = R_{u_{i+1}}(T) - R_{u_i}(T).$$

Since the variables $X_i(t)$'s are non-negative, we have $\sum_{i=1}^{k-1} X_{u_i}(t) \leq \sum_{j\neq r} X_j(t)$. Thus, for each node j

$$\frac{1}{T}\sum_{t=0}^{T-1} A(t) - \frac{1}{T}\sum_{j\neq r} X_j(T) \leq \frac{1}{T} R_j(T) \leq \frac{1}{T}\sum_{t=0}^{T-1} A(t).$$

Taking limit as $T \to \infty$ and using the strong law of large numbers for the arrival process and Eqn. (40), we have

$$\lim_{T\to\infty} \frac{R_j(T)}{T} = \lambda, \; \forall j. \quad \text{w.p. } 1$$

This concludes the proof.

8.3 Proof of Lemma (5)

We begin with a preliminary lemma.

Lemma 8. If we have

$$Q(t+1) \leq (Q(t) - \mu(t))^+ + A(t) \tag{42}$$

where all the variables are non-negative and $(x)^+ = \max\{x, 0\}$, then

$$Q^2(t+1) - Q^2(t) \leq \mu^2(t) + A^2(t) + 2Q(t)(A(t) - \mu(t)).$$

Proof: Squaring both sides of Eqn. (42) yields,

$$\begin{aligned} & Q^2(t+1) \\ & \leq \left((Q(t) - \mu(t))^+\right)^2 + A^2(t) + 2A(t)(Q(t) - \mu(t))^+ \\ & \leq (Q(t) - \mu(t))^2 + A^2(t) + 2A(t)Q(t), \end{aligned}$$

where we use the fact that $x^2 \geq (x^+)^2$, $Q(t) \geq 0$, and $\mu(t) \geq 0$. Rearranging the above inequality finishes the proof. □

Applying Lemma 8 to the dynamics (26) of $X_j(t)$ yields, for each node $j \neq r$,

$$X_j^2(t+1) - X_j^2(t) \leq B(t) + \tag{43}$$

$$2X_j(t)\Big(\sum_{m\in V} \mu_{mi^*_t}(t) - \sum_{k\in V} \mu_{kj}(t)\Big), \tag{44}$$

where $B(t) \leq c^2_{\max} + \max\{A^2(t), c^2_{\max}\} \leq (A^2(t) + 2c^2_{\max})$, $A(t)$ is the number of exogenous packet arrivals in a slot, and $c_{\max} \triangleq \max_{e\in E} c_e$ is the maximum capacity of the links. Since per-slot arrival $A(t)$ has finite second moment, there exists a finite constant $B > 0$ such that $\mathbb{E}[B(t)] \leq \mathbb{E}(A^2(t)) + 2c^2_{\max} < B$.

We define the quadratic Lyapunov function $L(\boldsymbol{X}(t)) = \sum_{j\neq r} X_j^2(t)$. From (43), the one-slot Lyapunov drift

$\Delta(\boldsymbol{X}(t))$, *conditioned* on the current network-configuration $\boldsymbol{\sigma}(t)$ yields

$$\begin{aligned}
& \Delta(\boldsymbol{X}(t)|\boldsymbol{\sigma}(t)) \triangleq \mathbb{E}[L(\boldsymbol{X}(t+1) - L(\boldsymbol{X}(t)) \mid \boldsymbol{X}(t), \boldsymbol{\sigma}(t)] \\
= \;& \mathbb{E}\big[\sum_{j\neq r}\big(X_j^2(t+1) - X_j^2(t)\big) \mid \boldsymbol{X}(t), \boldsymbol{\sigma}(t)\big] \\
\leq \;& B|V| + 2\sum_{j\neq r} X_j(t)\mathbb{E}\big[\sum_{m\in V}\mu_{mi_t^*}(t) \qquad (45)\\
& -\sum_{k\in V}\mu_{kj}(t) \mid \boldsymbol{X}(t), \boldsymbol{\sigma}(t)\big] \\
= \;& B|V| - 2\sum_{(i,j)\in E}\mathbb{E}[\mu_{ij}(t) \mid \boldsymbol{X}(t), \boldsymbol{\sigma}(t)]\big(X_j(t) \qquad (46)\\
& -\sum_{k\in K_j(t)} X_k(t)\big) \\
= \;& B|V| - 2\sum_{(i,j)\in E} W_{ij}(t)\mathbb{E}[\mu_{ij}(t) \mid \boldsymbol{X}(t), \boldsymbol{\sigma}(t)] \qquad (47)
\end{aligned}$$

The broadcast-policy π^* is chosen to minimize the upper-bound of *conditional-drift*, given on the right-hand side of (47) among all policies in Π^*.

Next, we construct a randomized scheduling policy $\pi^{\text{RAND}} \in \Pi^*$. Let $\boldsymbol{\beta}_{\boldsymbol{\sigma}}^* \in \text{conv}(\mathcal{M}_\sigma)$ be the part of an optimal solution corresponding to $\boldsymbol{\sigma}(t) \equiv \boldsymbol{\sigma}$ given by Eqn. 11. From Caratheodory's theorem [15], there exist at most $(|E|+1)$ link-activation vectors $\boldsymbol{s}_k \in \mathcal{M}_\sigma$ and the associated non-negative scalars $\{\alpha_k^\sigma\}$ with $\sum_{k=1}^{|E|+1}\alpha_k^\sigma = 1$, such that

$$\boldsymbol{\beta}_\sigma^* = \sum_{k=1}^{|E|+1}\alpha_k^\sigma \boldsymbol{s}_k^\sigma. \qquad (48)$$

Define the average (unconditional) activation vector

$$\boldsymbol{\beta}^* = \sum_{\sigma\in\Xi} p(\sigma)\boldsymbol{\beta}_\sigma^* \qquad (49)$$

Hence, from Eqn. (11) we have,

$$\lambda^* \leq \min_{U:\text{ a proper cut}}\sum_{e\in E_U} c_e\beta_e^*. \qquad (50)$$

Suppose that the exogenous packet arrival rate λ is strictly less than the broadcast capacity λ^*. There exists an $\epsilon > 0$ such that $\lambda + \epsilon \leq \lambda^*$. From (50), we have

$$\lambda + \epsilon \leq \min_{U:\text{ a proper cut}}\sum_{e\in E_U} c_e\beta_e^*. \qquad (51)$$

For any network node $v \neq r$, consider the proper cuts $U_v = V\setminus\{v\}$. Specializing the bound in (51) to these cuts, we have

$$\lambda + \epsilon \leq \sum_{e\in E_{U_v}} c_e\beta_e^*, \; \forall v \neq r. \qquad (52)$$

Since the underlying network topology $\mathcal{G} = (V, E)$ is a DAG, there exists a topological ordering of the network nodes so that: (i) the nodes can be labelled serially as $\{v_1, \ldots, v_{|V|}\}$, where $v_1 = r$ is the source node with no in-neighbours and $v_{|V|}$ has no outgoing neighbours and (ii) all edges in E are directed from $v_i \to v_j$, $i < j$ [4]; From (52), we define $q_l \in [0,1]$ for each node v_l such that

$$q_l\sum_{e\in E_{U_{v_l}}} c_e\beta_e^* = \lambda + \epsilon\frac{l}{|V|}, \; l = 2, \ldots, |V|. \qquad (53)$$

Consider the randomized broadcast policy $\pi^{\text{RAND}} \in \Pi^*$ working as follows:

Stationary Randomized Policy π^{RAND}:
(i) If the observed network-configuration at slot t is $\boldsymbol{\sigma}(t) = \sigma$, the policy π^{RAND} **selects** [2] the feasible activation set $\boldsymbol{s}_k^\sigma$ with probability α_k^σ;
(ii) For each incoming selected link $e = (\cdot, v_l)$ to node v_l such that $s_e(t) = 1$, the link e is **activated** independently with probability q_l;
(iii) **Activated** links (note, not necessarily all the *selected* links) are used to forward packets, subject to the constraints that define the policy class Π^* (i.e., in-order packet delivery and that a network node is only allowed to receive packets that have been received by all of its in-neighbors).

Note that this stationary randomized policy π^{RAND} operates independently of the state of received packets in the network, i.e., $\boldsymbol{X}(t)$. However it depends on the current network-configuration $\boldsymbol{\sigma}(t)$. Since each network node j is relabelled as v_l for some l, from (53) we have, for each node $j \neq r$, the total expected incoming transmission rate to the node j under the policy π^{RAND}, averaged over all network states $\boldsymbol{\sigma}$ satisfies

$$\begin{aligned}
\sum_{i:(i,j)\in E}\mathbb{E}[\mu_{ij}^{\pi^{\text{RAND}}}(t) \mid \boldsymbol{X}(t)] &= \sum_{i:(i,j)\in E}\mathbb{E}[\mu_{ij}^{\pi^{\text{RAND}}}(t)] \\
&= q_l\sum_{e\in E_{U_{v_l}}} c_e\beta_e^* \\
&= \lambda + \epsilon\frac{l}{|V|}. \qquad (54)
\end{aligned}$$

Equation (54) shows that the randomized policy π^{RAND} provides each network node $j \neq r$ with the total expected incoming rate strictly larger than the packet arrival rate λ via proper random link activations conditioned on the current network configuration. According to our notational convention, we have

$$\sum_{i:(i,r)\in E}\mathbb{E}[\mu_{ir}^{\pi^{\text{RAND}}}(t) \mid \boldsymbol{X}(t)] = \mathbb{E}[\sum_{i:(i,r)\in E}\mu_{ir}^{\pi^{\text{RAND}}}(t)] = \lambda. \qquad (55)$$

From (54) and (55), if node i appears before node j in the aforementioned topological ordering, i.e., $i = v_{l_i} < v_{l_j} = j$ for some $l_i < l_j$, then

$$\begin{aligned}
&\sum_{k:(k,i)\in E}\mathbb{E}[\mu_{ki}^{\pi^{\text{RAND}}}(t)] - \sum_{k:(k,j)\in E}\mathbb{E}[\mu_{kj}^{\pi^{\text{RAND}}}(t)] \\
&\leq -\frac{\epsilon}{|V|}. \qquad (56)
\end{aligned}$$

The above inequality will be used to show the throughput optimality of the policy π^*.
The drift inequality (45) holds for any policy $\pi \in \Pi^*$. The broadcast policy π^* observes the states $(\boldsymbol{X}(t), \boldsymbol{\sigma}(t))$ and and seek to *greedily* minimize the upper-bound of drift (47) at every slot. Comparing the actions taken by the policy π^* with those by the randomized policy π^{RAND} in slot t in (45),

we have

$$\Delta^{\pi^*}(\boldsymbol{X}(t)|\boldsymbol{\sigma}(t)) \tag{57}$$
$$\leq B|V| - 2\sum_{(i,j)\in E}\mathbb{E}[\mu_{ij}^{\pi^*}(t) \mid \boldsymbol{X}(t), \boldsymbol{\sigma}(t)]W_{ij}(t)$$
$$\leq B|V| - 2\sum_{(i,j)\in E}\mathbb{E}[\mu_{ij}^{\pi^{\text{RAND}}}(t) \mid \boldsymbol{X}(t), \boldsymbol{\sigma}(t)]W_{ij}(t)$$
$$\overset{(*)}{=} B|V| - 2\sum_{(i,j)\in E}\mathbb{E}[\mu_{ij}^{\pi^{\text{RAND}}}(t) \mid \boldsymbol{\sigma}(t)]W_{ij}(t) \tag{58}$$

Taking Expectation of both sides w.r.t. the stationary-process $\boldsymbol{\sigma}(t)$ and rearranging, we have

$$\Delta^{\pi^*}(\boldsymbol{X}(t)) \tag{59}$$
$$\leq B|V| - 2\sum_{(i,j)\in E}\mathbb{E}[\mu_{ij}^{\pi^{\text{RAND}}}(t)]W_{ij}(t)$$
$$\leq B|V| + 2\sum_{j\neq r}X_j(t)\Big(\sum_{m\in V}\mathbb{E}[\mu_{mi_t^*}^{\pi^{\text{RAND}}}(t)] - \sum_{k\in V}\mathbb{E}[\mu_{kj}^{\pi^{\text{RAND}}}(t)]\Big)$$
$$\leq B|V| - \frac{2\epsilon}{|V|}\sum_{j\neq r}X_j(t). \tag{60}$$

Note that $i_t^* = \arg\min_{i\in \text{In}(j)} Q_{ij}(t)$ for a given node j. Since node i_t^* is an in-neighbour of node j, i_t^* must lie before j in any topological ordering of the DAG. Hence, the last inequality of (60) follows directly from (56). Taking expectation in (60) with respect to $\boldsymbol{X}(t)$, we have

$$\mathbb{E}\big[L(\boldsymbol{X}(t+1))\big] - \mathbb{E}\big[L(\boldsymbol{X}(t))\big] \leq B|V| - \frac{2\epsilon}{|V|}\mathbb{E}||\boldsymbol{X}(t)||_1,$$

where $||\cdot||_1$ is the ℓ_1-norm of a vector. Summing the above inequality over $t = 0, 1, 2, \ldots T-1$ yields

$$\mathbb{E}\big[L(\boldsymbol{X}(T))\big] - \mathbb{E}\big[L(\boldsymbol{X}(0))\big] \leq B|V|T - \frac{2\epsilon}{|V|}\sum_{t=0}^{T-1}\mathbb{E}||\boldsymbol{X}(t)||_1.$$

Dividing the above by $2T\epsilon/|V|$ and using $L(\boldsymbol{X}(t)) \geq 0$, we have

$$\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}||\boldsymbol{X}(t)||_1 \leq \frac{B|V|^2}{2\epsilon} + \frac{|V|\,\mathbb{E}[L(\boldsymbol{X}(0))]}{2T\epsilon}$$

Taking a $\limsup$ of both sides yields

$$\limsup_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\sum_{j\neq r}\mathbb{E}[X_j(t)] \leq \frac{B|V|^2}{2\epsilon} \tag{61}$$

which implies that all virtual-queues $X_j(t)$ are strongly stable [16]. Strong stability of $X_j(t)$ implies that all virtual queues $X_j(t)$ are rate stable [16, Theorem 2.8].

8.4 Proof of Lemma (6)

Proof: Recall the definition of the policy-space Π^*. For every node i, since $R_i(t)$ is a *non-decreasing function* of t, if a packet p is allowed to be transmitted to a node j at time slot t, by the policy π', it is certainly allowed to be transmitted by the policy π. This is because $R'_i(t) \leq R_i(t), \forall j \in \partial^{\text{out}}(i)$ and hence outdated state-information may only prevent transmission of a packet p at a time t, which would otherwise be allowed by the policy π^*. As a result, the policy π' can never transmit a packet to node j which is not present at all in-neighbours of the node j. This shows that $\pi' \in \Pi^*$. □

8.5 Proof of Lemma (7)

Consider the packet-state update process at node j. Since the capacity of the links are bounded by $c_{\max}$, from Eqn. (31) and the fact that $R_i(t)$ is non-decreasing, we have

$$R_i(t) - Tc_{\max} \leq R'_i(t) \leq R_i(t), \ \ \forall i \in \partial^{\text{in}}(j) \tag{62}$$

Hence, from Eq. (24), it follows that

$$X_j(t) - Tc_{\max} \leq X'_j(t) \leq X_j(t) \tag{63}$$

i.e.,

$$X_j(t) - c_{\max}\mathbb{E}T \leq \mathbb{E}X'_j(t) \leq X_j(t) \tag{64}$$

Where the expectation is with respect to the random update process at the node j. In a similar fashion, since every in-neighbour i of a node $k \in \partial^{\text{out}}(j)$, is at most 2-hop away from the node j, we have

$$R_i(t) - Tc_{\max} \leq R'_i(t) \leq R_i(t)$$

Also,

$$R_k(t) - Tc_{\max} \leq R'_k(t) \leq R_k(t)$$

It follows that for all $i \in \partial^{\text{in}}(k)$

$$\begin{aligned}(R_i(t) - R_k(t)) - Tc_{\max} &\leq R'_i(t) - R'_k(t)\\ &\leq (R_i(t) - R_k(t)) + Tc_{\max}\end{aligned}$$

Hence,

$$X_k(t) - Tc_{\max} \leq X'_k(t) \leq X_k(t) + Tc_{\max}$$

Again taking expectation w.r.t. the random packet-state update process,

$$X_k(t) - c_{\max}\mathbb{E}T \leq \mathbb{E}X'_k(t) \leq X_k(t) + c_{\max}\mathbb{E}T \tag{65}$$

Combining Eqns (64) and (65) using Linearity of expectation and using Eqn. (28) we have

$$-nc_{\max}\mathbb{E}T + W_{ij}(t) \leq \mathbb{E}W'_{ij}(t) \leq W_{ij}(t) + nc_{\max}\mathbb{E}T$$

Thus the lemma (7) follows with $C \equiv nc_{\max}\mathbb{E}T < \infty$.

8.6 Proof of Theorem (5.1)

To prove throughput-optimality of Theorem (5.1), we work with the same Lyapunov function $L(\boldsymbol{X}(t)) = \sum_{j\neq r}X_j^2(t)$ as in Theorem (4.1) and follow the same steps until Eqn. (47) to obtain the following upper-bound on conditional drift

$$\begin{aligned}&\Delta^{\pi'}(\boldsymbol{X}(t)|\boldsymbol{X}(t), \boldsymbol{X}'(t), \boldsymbol{\sigma}(t))\\ \leq\ & B|V| - 2\sum_{(i,j)\in E}W_{ij}(t)\mathbb{E}(\mu_{ij}^{\pi'}(t)|\boldsymbol{X}(t), \boldsymbol{X}'(t), \boldsymbol{\sigma}(t))\end{aligned} \tag{66}$$

Since the policy π' makes scheduling decision based on the *locally computed* weights $W'_{ij}(t)$, by the definition of the policy π', we have for any policy $\pi \in \Pi$:

$$\begin{aligned}&\sum_{(i,j)\in E}W'_{ij}(t)\mathbb{E}(\mu_{ij}^{\pi'}(t)|\boldsymbol{X}(t), \boldsymbol{X}'(t), \boldsymbol{\sigma}(t))\\ \geq &\sum_{(i,j)\in E}W'_{ij}(t)\mathbb{E}(\mu_{ij}^{\pi}(t)|\boldsymbol{X}(t), \boldsymbol{X}'(t), \boldsymbol{\sigma}(t))\end{aligned} \tag{67}$$

Taking expectation of both sides w.r.t. the random update process $\boldsymbol{X}'(t)$, conditioned on the true network state $\boldsymbol{X}(t)$ and the network configuration $\boldsymbol{\sigma}(t)$, we have

$$\begin{aligned}
&Cn^2c_{\max}/2 + \sum_{(i,j)\in E} W_{ij}(t)\mathbb{E}(\mu_{ij}^{\pi'}(t)|\boldsymbol{X}(t),\boldsymbol{\sigma}(t))\\
\overset{(a)}{\geq}\ & \sum_{(i,j)\in E} \mathbb{E}W'_{ij}(t)\mathbb{E}(\mu_{ij}^{\pi'}(t)|\boldsymbol{X}(t),\boldsymbol{\sigma}(t))\\
\overset{(b)}{\geq}\ & \sum_{(i,j)\in E} \mathbb{E}W'_{ij}(t)\mathbb{E}(\mu_{ij}^{\pi}(t)|\boldsymbol{X}(t),\boldsymbol{\sigma}(t))\\
\overset{(c)}{\geq}\ & \sum_{(i,j)\in E} W_{ij}(t)\mathbb{E}(\mu_{ij}^{\pi}(t)|\boldsymbol{X}(t),\boldsymbol{\sigma}(t)) - Cn^2c_{\max}/2
\end{aligned} \tag{68}$$

Here the inequality (a) and (c) follows from Lemma (7) and the fact that $|E| \leq n^2/2$ and $\mu_{ij}(t) \leq c_{\max}$. Inequality (b) follows from Eqn. (67). Thus from Eqn. (66) and (68), the expected conditional drift of the Lyapunov function under the policy π', where the expectation is taken w.r.t. the random update and arrival process is upper-bounded as follows:

$$\Delta^{\pi'}(\boldsymbol{X}(t)|\boldsymbol{X}(t),\boldsymbol{\sigma}(t)) \leq B' - 2\sum_{(i,j)\in E} W_{ij}(t)\mathbb{E}[\mu_{ij}^{\pi}(t) \mid \boldsymbol{X}(t),$$

with the constant $B' \equiv B|V| + 2Cn^2c_{\max}$. Since the above inequality holds for any policy $\pi \in \Pi$, we can follow the exactly same steps in the proof of Theorem (4.1) by replacing an arbitrary π by π^{RAND} and showing that it has negative drift.

8.7 Proof of Proposition 6.1

Like many proofs in this paper, this proof also has a converse and an achievability part. In the converse part, we obtain an upper bound of $\frac{2}{5}$ for the broadcast capacity λ^*_{stat} of the stationary grid network (i.e. when all links are ON w.p. 1). In the achievability part, we show that this upper bound is tight.

Part (a): Proof of the Converse: $\lambda^*_{\text{stat}} \leq \frac{2}{5}$

We have shown earlier that for the purpose of achieving capacity, it is sufficient to restrict our attention to stationary randomized policies only. Suppose a stationary randomized policy π achieves a broadcast rate λ and it activates edge $e \in E$ at every slot with probability f_e. Then for the nodes $\mathtt{a}$ and $\mathtt{b}$ to receive distinct packets at rate λ, one requires

$$f_{\mathtt{ra}} \geq \lambda, f_{\mathtt{ab}} \geq \lambda$$

Applying the primary interference constraint at node $\mathtt{a}$, we then obtain

$$f_{\mathtt{ad}} \leq 1 - 2\lambda$$

Because of symmetry in the network topology, we also have

$$f_{\mathtt{cd}} \leq 1 - 2\lambda.$$

However, to achieve a broadcast capacity of λ, the total allocated rate towards node $\mathtt{d}$ must be atleast λ. Hence,

$$2(1-2\lambda) \geq \lambda$$

i.e.

$$\lambda \leq \frac{2}{5}.$$

Since the above holds for any stationary randomized policy π, we conclude

$$\lambda^*_{\text{stat}} \leq \frac{2}{5} \tag{69}$$

■

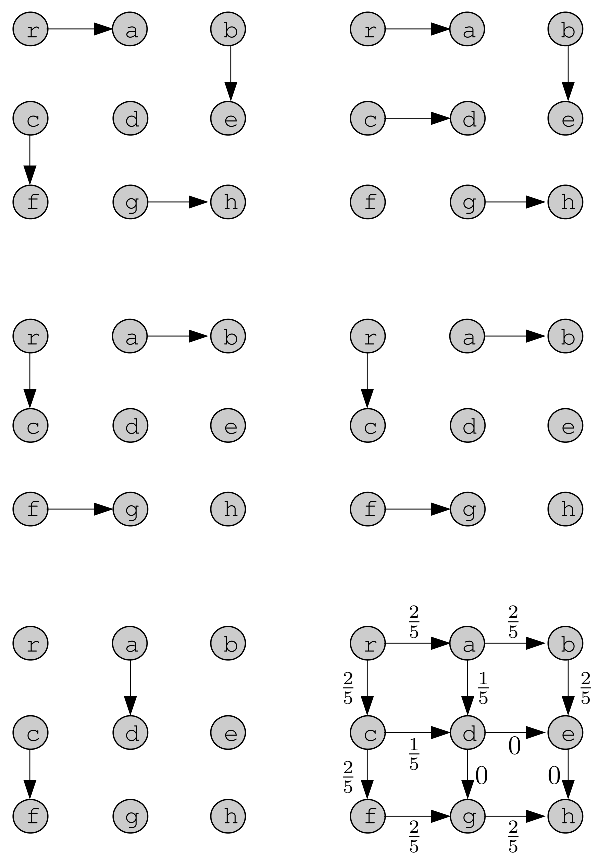

Matching M_5 'Time averaged' Network

Fig. 6. Some feasible activations of the 3×3 grid network which are activated uniformly at random. The components corresponding to each edge in the resulting overall activation vector $\boldsymbol{\beta}$ is denoted by the numbers alongside the edges.

Part (b): Proof of the Achievability: $\lambda^*_{\text{stat}} \geq \frac{2}{5}$*:*

As usual, the achievability proof will be constructive. Consider the following five activations (matchings) $M_1, M_2, \ldots, M_5$ of the underlying graph as shown in Figure 6. Now consider a stationary policy $\pi^* \in \Pi^*$ that activates the matchings $M_1, \ldots, M_5$ at each slot uniformly at random with probability $\frac{1}{5}$ for each matching. The resulting 'time-averaged' graph is also shown in Figure 6. Using Theorem 3.1, it is clear that $\lambda^*_{\text{stat}} \geq \frac{2}{5}$. Combining the above with the converse result in Eqn. (69), we conclude that,

$$\lambda^*_{\text{stat}} = \frac{2}{5}$$

■