

On the Power of Randomization for Scheduling Real-Time Traffic in Wireless Networks

Christos Tsanikidis, Javad Ghaderi
Electrical Engineering Department, Columbia University

Abstract—In this paper, we consider the problem of scheduling real-time traffic in wireless networks under a conflict-graph interference model and single-hop traffic. The objective is to guarantee that at least a certain fraction of packets of each link are delivered within their deadlines, which is referred to as *delivery ratio*. This problem has been studied before under restrictive frame-based traffic models, or greedy maximal scheduling schemes like LDF (Largest-Deficit First) that can lead to poor delivery ratio for general traffic patterns. In this paper, we pursue a different approach through randomization over the choice of maximal links that can transmit at each time. We design randomized policies in collocated networks, multi-partite networks, and general networks, that can achieve delivery ratios much higher than what is achievable by LDF. Further, our results apply to traffic (arrival and deadline) processes that evolve as positive recurrent Markov chains. Hence, this work is an improvement with respect to both efficiency and traffic assumptions compared to the past work. We further present extensive simulation results over various traffic patterns and interference graphs to illustrate the gains of our randomized policies over LDF variants.

Index Terms—Scheduling, Real-Time Traffic, Markov Processes, Stability, Wireless Networks

I. INTRODUCTION

Much of the prior work on scheduling algorithms for wireless networks focus on maximizing throughput. However, for many real-time applications, e.g., in Internet of Things (IoT), vehicular networks, and other cyber-physical systems, delays and deadline guarantees on packet delivery are more important than long-term throughput [1]–[3]. Recently, there has been an interest in developing scheduling algorithms specifically targeted towards handling deadline-constrained traffic [4]–[9], when each packet has to be delivered within a *strict deadline*, otherwise it is of no use. The key objective in these works is to guarantee that at least a fraction of the packets will be delivered to their destinations within their deadlines, which is referred to as *delivery ratio* (QoS). Providing such guarantees is very challenging as it crucially depends on the temporal pattern of packet arrivals and their deadlines, as opposed to long-term averages in traditional throughput maximization. One can construct adversarial traffic patterns that all have the same long-term average but their achievable delivery ratio is vastly different [8], [10].

Recently, there have been two approaches for providing QoS guarantees for real-time traffic in wireless networks.

Emails: {c.tsanikidis, jghaderi}@columbia.edu. This research was supported by grants NSF 1717867, NSF 1652115, and ARO W911NF1910379.

One is the frame-based approach [4]–[7], and the other is a greedy scheduling approach like the largest-deficit-first policy (LDF) [8], [9]. In the frame-based approach, it is assumed that each frame is a number of consecutive time slots, and packets arriving in each frame have to be scheduled before the end of the frame. They crucially rely on the assumption that all packets of all users arrive at the beginning of frames [4]–[6], or the complete knowledge of future packet arrivals and their deadlines in each frame is available at the beginning of the frame [7]. This restricts the application of such policies to specific traffic patterns with periodic arrivals and synchronized users. The results for general traffic patterns without such frame assumptions are very limited, as in such settings, the real-time rate region is difficult to characterize and the optimal policy is unknown. A popular algorithm for providing QoS guarantees for real-time traffic is the largest-deficit-first (LDF) policy [4], [8], [9], [11], which is the real-time variation of the longest-queue-first (LQF) policy (see, e.g., [12], [13]). It is known that LDF is optimal in collocated networks *under the frame-based model* [4], [11]. The performance of LDF in the non-frame-based setting has been studied in [8] in terms of the *efficiency ratio*, which is the fraction of the real-time throughput region guaranteed by LDF. It is shown that LDF achieves an efficiency ratio of at least $\frac{1}{1+\beta}$ for a network with interference degree¹ β , under *i.i.d.* (independent and identically distributed) packet arrivals and deadlines. Further, when traffic is not *i.i.d.*, the efficiency ratio of LDF is as low as $\frac{1}{1+\sqrt{\beta}}$ [8]. In particular, for collocated networks, the efficiency ratio of LDF under non-*i.i.d.* traffic is $1/2$, and in a simple star topology with one center link and K neighboring links, it scales down as low as $O(\frac{1}{\sqrt{K}})$. This shows that LDF might not be suitable for high throughput real-time applications, especially with non-*i.i.d.* traffic, which is the case if packet drops due to deadline expiry trigger re-transmissions.

Besides the works above on providing QoS guarantees for wireless networks, there is literature on approximation algorithms for *single-link* buffer management problem [14], [15]. In this problem, packets arrive to a single link, each with a non-negative constant weight and a deadline. The goal is to maximize the total weight of transmitted packets for the worst input sequence. The approximation algorithms include the maximum-weight greedy algorithm [14], [15], EDF _{α} [16]

¹The interference degree is the maximum number of links that can be scheduled simultaneously out of a link and its neighboring links.

which schedules the earliest-deadline packet with weight at least $\alpha \leq 1$ of the maximum-weight packet, or randomized algorithms such as [17]–[19] where the scheduling decision is randomized over pending packets in the link’s buffer. Inspired by such randomization techniques, we design randomized algorithms for wireless networks under a general interference model and given the delivery ratio requirements for the links in the network.

A. Contributions

Non-i.i.d. (Markovian) Traffic Models. Our traffic model allows traffic (*arrival* and *deadline*) processes that evolve as an irreducible Markov chain over a finite state space. This model is a significant extension from i.i.d. or frame-based traffic models in [4]–[8]. A key technique in analyzing the achievable efficiency ratio in our model is to look at the return times of the traffic Markov chain and analyze the performance of scheduling algorithms over long enough cycles consisting of multiple return times.

Randomized Algorithms with Improved Efficiency. We propose randomized scheduling algorithms that can significantly outperform deterministic greedy algorithms like LDF. The key idea is to identify a structure for the optimal policy and randomize over the possible scheduling choices of the optimal policy, rather than solely relying on the deficit queues. For *collocated networks* and *complete bipartite graphs* our randomized algorithms achieve an efficiency ratio of at least 0.63 and 2/3, respectively, and in *general graphs*, achieve an efficiency ratio of at least 1/2, all *independent of the network size and without the knowledge of the traffic model*.

II. MODEL AND DEFINITIONS

Wireless Network Model. We consider a set of K links (or users) denoted by the set \mathcal{K} , where $K = |\mathcal{K}|$. Time is slotted, and at each time slot $t \in \mathbb{N}_0$, each link can transmit one packet successfully, if there are no interfering links transmitting at the same time. As in [8], it is standard to represent the interference relationships between links by an *interference graph* $G_I = (\mathcal{K}, E_I)$. Each vertex of G_I is a link, and an edge $(l_1, l_2) \in E_I$ indicates links l_1 and l_2 interfere with each other. Let $I_l(t) = 1$ if link l is transmitting a packet at time t , and $I_l(t) = 0$ otherwise. Hence, at any time any feasible schedule $I(t) = (I_l(t), l \in \mathcal{K})$ has to form an independent set of G_I over links that have packets, i.e., no two transmitting links can share an edge in G_I . We say a feasible schedule I is maximal if no more links can be scheduled without interfering with some active links in I . Let $\mathcal{B}(t)$ be the set of links that have packets available to transmit at time t . Let \mathcal{M} denote the set of all maximal independent sets of G_I . Then, at any time t ,

$$\{l \in \mathcal{K} : I_l(t) = 1\} \subseteq (\mathcal{B}(t) \cap M), \text{ for some } M \in \mathcal{M},$$

where ‘ \subseteq ’ holds with ‘ $=$ ’ if I is a maximal schedule.

Traffic Model. We consider a single-hop traffic with deadlines for each link. Let $a_l(t)$ denote the number of packets arriving on link l at time t , with $a_l(t) \leq a_{\max}$, for some $a_{\max} < \infty$. Each packet upon arrival has a deadline which is




Fig. 1: An example with three traffic patterns $A \rightarrow B \rightarrow C \rightarrow \dots$. Each rectangle indicates a packet for a link indicated by its number. The left side of the rectangle corresponds to its arrival time, and the right side corresponds to its deadline.

the maximum delay that the packet can tolerate. We define a vector $\tau_l(t) = (\tau_{l,d}(t); d = 1, \dots, d_{\max})$, where $\tau_{l,d}(t)$ is the number of packets with deadline d arriving to link l at time t . A packet arriving with deadline d at time t has to be transmitted before the end of time slot $t + d - 1$, otherwise it will be dropped. The maximum deadline is bounded by a constant d_{\max} . Hence, the network traffic (arrival and deadline) process is described by $\tau(t) = (\tau_l(t); l \in \mathcal{K})$, $t \geq 0$. We assume that the traffic process $\tau(t)$ evolves as an (unknown) irreducible Markov chain over the space $\mathcal{X} = \{0, \dots, a_{\max}\}^{d_{\max} \times K}$ ².

Note that the arrival and deadline processes do not need to be i.i.d. across times or users. Since the state space \mathcal{X} is finite, $\tau(t)$ is a positive recurrent Markov chain [21] and the time-average of any bounded function of $\tau(t)$ is well-defined, in particular, the packet arrival rate \bar{a}_l , $l \in \mathcal{K}$,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t a_l(s) = \bar{a}_l. \quad (1)$$

See Figure 1 for an example of a Markovian traffic process.

Buffer Dynamics. The buffer of link l at time t contains the existing packets at link l which have not expired yet and also the newly arrived packets $\tau_l(t)$. Formally, we define the buffer of link l by a vector $\Psi_l(t) = (\Psi_{l,d}(t); d = 1, \dots, d_{\max})$, where $\Psi_{l,d}(t)$ is the number of packets in the buffer with *remaining deadline* d at time t . The remaining deadline of each packet in the buffer decreases by one at every time slot, until the packet is successfully transmitted or reaches the deadline 0, which in either case the packet is removed from the buffer, i.e., the buffer at the beginning of slot $t + 1$ is

$$\Psi_{l,d}(t+1) = \Psi_{l,d+1}(t) + \tau_{l,d}(t+1) - I_{l,d+1}(t), \quad (2)$$

where $I_l(t) = \sum_{d=1}^{d_{\max}} I_{l,d}(t) \leq 1$, and $I_{l,d}(t) = 1$ if the scheduler selects a packet with deadline d to transmit at time t on link l . By convention, we set $\Psi_{l,d_{\max}+1}(t) = 0$, $\Psi_{l,0}(t) = 0$. We define the network buffer state as $\Psi(t) = (\Psi_l(t); l \in \mathcal{K})$.

Delivery Requirement and Deficit. As in [4]–[8], we assume that there is a minimum delivery ratio p_l (QoS requirement) for each link l , $l \in \mathcal{K}$. This means the scheduling algorithm must successfully deliver at least p_l fraction of the incoming packets on each link l in long term. Formally,

$$\liminf_{t \rightarrow \infty} \frac{\sum_{s=1}^t I_l(s)}{\sum_{s=1}^t a_l(s)} \geq p_l. \quad (3)$$

²The state can be extended to $(\tau(u), u(t))$ where $u(t)$ is the unobservable part which can encode more complicated finite state Markov chains for the traffic process (see our technical report [20]).

We define a deficit $w_l(t)$ which measures the amount of service owed to link l up to time t to fulfill its minimum delivery rate. As in [7], [8], the deficit evolves as

$$w_l(t+1) = [w_l(t) + \tilde{a}_l(t) - I_l(t)]^+, \quad (4)$$

where $[\cdot]^+ = \max\{\cdot, 0\}$, and $\tilde{a}_l(t)$ indicates the amount of deficit increase due to packet arrivals. For each packet arrival, we should increase the deficit by p_l on average. For example, we can increase the deficit by exactly p_l for each packet arrival to link l , or use a coin tossing process as in [7], [8], i.e., each packet arrival at link l increases the deficit by one with the probability p_l , and zero otherwise. We refer to $\tilde{a}_l(t)$ as the *deficit arrival* process for link l . Note that it holds that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \tilde{a}_l(s) = \bar{a}_l p_l := \lambda_l, \quad l \in \mathcal{K}. \quad (5)$$

We refer to λ_l as the deficit arrival rate for link l . We emphasize that the arriving packet is always added to the link's buffer, regardless of whether and how much deficit is added for that packet. Also note that in (4), each time a packet is scheduled from the link, $I_l(t) = 1$, the deficit is reduced by one. The dynamics in (4) define a deficit queueing system, with bounded increments/decrements, whose stability, e.g.,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \mathbb{E}[w_l(s)] < \infty, \quad (6)$$

implies that (3) holds³. Define the vector of deficits as $w(t) = (w_l(t), l \in \mathcal{K})$. The system state at time t is then defined as $\mathcal{S}(t) = (\Psi(t), w(t), \tau(t))$.

Objective. Define \mathcal{P}_C to be the set of all causal policies, i.e. policies that do not know the information of future arrivals and deadlines in order to make scheduling decisions. For a given traffic process $\tau(t)$, with fixed \bar{a}_l , defined in (1), we are interested in causal policies that can stabilize the deficit queues for the largest set of delivery rate vectors $\mathbf{p} = (p_l, l \in \mathcal{K})$, or equivalently largest set of $\lambda = (\lambda_l := \bar{a}_l p_l, l \in \mathcal{K})$ possible. For a given traffic process, we say the rate vector $\lambda = (\lambda_l, l \in \mathcal{K})$ is supportable under some policy $\mu \in \mathcal{P}_C$ if all the deficit queues remain stable. Then one can define the supportable (real-time) rate region of the policy μ as

$$\Lambda_\mu = \{\lambda \geq 0 : \lambda \text{ is supportable by } \mu\}. \quad (7)$$

Note that for a given traffic distribution, a vector λ corresponds to a single vector of delivery rate requirements \mathbf{p} exactly. The supportable rate region under all the causal policies is defined as $\Lambda = \bigcup_{\mu \in \mathcal{P}_C} \Lambda_\mu$. The overall performance of a policy μ is evaluated by the efficiency ratio γ_μ^* which is defined as

$$\gamma_\mu^* = \sup\{\gamma : \gamma \Lambda \subseteq \Lambda_\mu\}. \quad (8)$$

For a causal policy μ , we aim to provide a *universal lower bound* on the efficiency ratio that holds for “all” Markovian traffic processes.

³Actually only the rate stability is enough to establish (3) [22], however we consider this stronger notion of stability.




Fig. 2: An example for non-dominated links. Each numbered rectangle denotes the earliest-deadline packet of a link. A solid rectangle indicates that the link is non-dominated. Dashed rectangles (links) that fall in regions \mathcal{R}_i will be dominated.

III. RANDOMIZED SCHEDULING ALGORITHMS

In this section, we present our randomized scheduling algorithms. We start with the collocated networks, and then proceed to general networks.

A. Collocated Networks

In a collocated network, only one of the links can transmit a packet at any time. Hence the interference graph G_I is a complete graph.

Define $e_l(t) = \min\{d : \Psi_{l,d}(t) > 0\}$ to be the deadline of the earliest-deadline packet available at link l at time t . By convention, the minimum of an empty set is considered infinity. We use a tuple $(w_l(t), e_l(t))_l$ to denote the earliest-deadline packet of link l with deadline $e_l(t)$ and link deficit $w_l(t)$. We make the following dominance definition.

Definition 1. We say that a link l_1 dominates a link l_2 at time t if $w_{l_1}(t) \geq w_{l_2}(t)$ and $e_{l_1}(t) \leq e_{l_2}(t)$. If one of the two inequalities is strict, we call it a strict dominance. A non-dominated link is a nonempty link that is not dominated strictly by any other link at that time.

Recall that $\mathcal{B}(t)$ is the set of links with nonempty buffers. At every time slot, we first find the set of non-dominated links $\mathcal{B}_{\text{ND}}(t)$. One way to do that is as follows:

Algorithm 1 Finding the Set Of Non-dominated Links

```

1:  $H \leftarrow \mathcal{B}(t)$ ,  $\mathcal{B}_{\text{ND}}(t) \leftarrow \emptyset$ ,  $i \leftarrow 0$ 
2: while  $H \neq \emptyset$  do
3:    $i \leftarrow i + 1$ 
4:   Find the largest-deficit non-dominated link  $h_i \in H$ .
5:   Add  $h_i$  to  $\mathcal{B}_{\text{ND}}(t)$ 
6:   Remove  $h_i$  and all the links dominated by it, i.e.
     $H \leftarrow H \setminus \{l \in H : e_l(t) \geq e_{h_i}(t)\}$ .
7: end while

```

Algorithm 1 returns a set $\mathcal{B}_{\text{ND}}(t) = \{h_1, \dots, h_k\}$, where h_i is the link selected in the i -th iteration, and the links are ordered in the order of their deficits, i.e., $w_{h_1}(t) > w_{h_2}(t) > \dots > w_{h_k}(t)$. See Figure 2 for an illustrative example of the non-dominated links. Our scheduling algorithm transmits the earliest-deadline packet of one of the links $h_i \in \mathcal{B}_{\text{ND}}(t)$ randomly, where the probabilities $p_{h_i}(t)$ are computed recursively

Algorithm 2 AMIX-ND: Randomized Scheduling in Collocated Networks

```

1: Use Algorithm 1 to find  $\mathcal{B}_{\text{ND}}(t) = \{h_1, \dots, h_k\}$ .
2:  $r \leftarrow 1$ 
3: for  $i = 1$  to  $k - 1$  do
4:    $p_{h_i}(t) = \min\left(1 - \frac{w_{h_{i+1}}(t)}{w_{h_i}(t)}, r\right)$ 
5:    $r \leftarrow r - p_{h_i}(t)$ 
6: end for
7:  $p_{h_k}(t) = r$ 
8: Send the earliest-deadline packet from link  $h_i$  with probability  $p_{h_i}(t)$ .

```

as in Algorithm 2. We refer to Algorithm 2 as AMIX-ND which stands for *Adaptive Mixing over Non-Dominated links*.

Theorem 1. *In a collocated wireless network with K links, AMIX-ND achieves an efficiency ratio of at least*

$$\gamma_{\text{AMIX-ND}}^* \geq 1 - \left(1 - \frac{1}{K}\right)^K > \frac{e-1}{e}. \quad (9)$$

Remark 1. Note that AMIX-ND has an efficiency ratio which is bounded from below by 0.63, regardless of the number of links. In contrast, we can construct Markovian traffic processes where the efficiency ratio of LDF is less than $1/2 + \epsilon$ [8]. For example, for the traffic patterns of Figure 1 in the model section, we will see in simulations in Section VI that, while AMIX-ND can achieve delivery ratios close to 0.99, LDF can not do better than $0.5 + \epsilon$. Note that our traffic model does allow traffic patterns as in Figure 1, since we do not need the traffic Markov chain to be aperiodic.

B. Multipartite Networks and General Networks

Consider the set of all maximal independent sets \mathcal{M} of the interference graph G_I . Our randomized algorithm selects a maximal independent set (MIS) $M \in \mathcal{M}$ probabilistically and schedules the earliest-deadline packets of the induced maximal schedule $M \cap \mathcal{B}(t)$. Recall that $\mathcal{B}(t)$ is the set of links with nonempty buffers. We refer to this algorithm as AMIX-MS which stands for *Adaptive Mixing over Maximal Schedules*. Before presenting the algorithm, we make a few definitions.

Definition 2. *The weight of a MIS $M \in \mathcal{M}$ at time t is*

$$W_M(t) = \sum_{l \in M \cap \mathcal{B}(t)} w_l(t). \quad (10)$$

Let $R = |\mathcal{M}|$. We index and order $M \in \mathcal{M}$ such that M_i has the i -th largest weight at time t , i.e.,

$$W_{M_1}(t) \geq W_{M_2}(t) \cdots \geq W_{M_R}(t).$$

Definition 3. *Define the subharmonic average of weights of the first n MIS, $n \leq R$, at time t to be*

$$C_n(t) = \frac{n-1}{\sum_{i=1}^n (W_{M_i}(t))^{-1}}. \quad (11)$$

The probabilities used by AMIX-MS to select a MIS M_i , at time t , are as follows

$$p_{M_i}^{\bar{n}}(t) \equiv p_i^{\bar{n}}(t) = \begin{cases} 1 - \frac{C_n(t)}{W_{M_i}(t)} & 1 \leq i \leq \bar{n} \\ 0 & \bar{n} < i \leq R \end{cases} \quad (12)$$

where \bar{n} is the largest n such that $\{p_i^n(t), 1 \leq i \leq n\}$ defines a valid probability distribution over $1 \leq i \leq n$. Noting that $p_i^n(t) \geq p_{i+1}^n(t)$ for $i < n$, and $\sum_{i \leq n} p_i^n(t) = 1$, \bar{n} is therefore given by

$$\bar{n} := \bar{n}(t) = \max\{n : p_n^n(t) \geq 0\}. \quad (13)$$

We drop the dependence on t for $\bar{n}(t)$ when there is no ambiguity. Algorithm 3 gives a description of AMIX-MS where \bar{n} is found using binary search. Then AMIX-MS selects a MIS M_i with probability $p_i^{\bar{n}}(t)$ as in (12).

Algorithm 3 AMIX-MS: Randomized Scheduling in General Interference Graphs

```

1:  $n_1 \leftarrow 1, n_2 \leftarrow |\mathcal{M}|$ 
2: while  $n_1 \neq n_2$  do
3:    $n \leftarrow \lceil \frac{n_1+n_2}{2} \rceil$ 
4:   if  $p_n^n(t) \geq 0$  then
5:      $n_1 \leftarrow n$ 
6:   else
7:      $n_2 \leftarrow n - 1$ 
8:   end if
9: end while
10:  $\bar{n} \leftarrow n_1$ 
11: Select MIS  $M_i$  with probability  $p_{M_i}^{\bar{n}}(t)$  as in (12) and transmit the earliest-deadline packet of each link in  $M_i$ .

```

The following theorem states the main result regarding the efficiency ratio of AMIX-MS.

Theorem 2. *In a wireless network with interference graph G_I and maximal independent sets \mathcal{M} , the efficiency ratio of AMIX-MS is at least*

$$\gamma_{\text{AMIX-MS}}^* \geq \frac{|\mathcal{M}|}{2|\mathcal{M}| - 1} > \frac{1}{2}.$$

A special case of this theorem is for networks with a complete n -partite interference graph, $n \geq 2$. In a complete n -partite graph, with n components, V_1, \dots, V_n , links in each component do not share any edge but there is an edge between any two links in different components. Hence, each component V_i , $1 \leq i \leq n$ is a MIS. We state the result as the following corollary which immediately follows from Theorem 2.

Corollary 2.1. *For a wireless network with a complete n -partite interference graph, under AMIX-MS,*

$$\gamma_{\text{AMIX-MS}}^* \geq \frac{n}{2n-1}.$$

Remark 2. We emphasize on the importance of Theorem 2 using a simple interference graph with ‘star’ topology. This is a special case of a bipartite graph with only two components,

V_1 is the center node, and V_2 are the leaf nodes. Notice that the guarantee of AMIX-MS in this case is at least $\frac{2}{3}$, regardless of the number of nodes K . This is a significant improvement over LDF, whose efficiency ratio is at least $\frac{1}{K}$ under i.i.d. traffic but not better than $\frac{1}{\sqrt{K-1}+1}$ under Markovian traffics [8].

Remark 3. We note that the computational complexity of AMIX-MS could be high for general graphs as it requires finding an ordering of maximal schedules. However, it is easily applicable for n -partite graphs or small graphs. Moreover, we can further approximate the algorithm by only ordering a subset of maximal schedules as opposed to finding all of them. The randomization in AMIX-MS can be also potentially implemented in a distributed manner by using distributed CSMA-like schemes such as [23]–[25].

IV. ANALYSIS TECHNIQUE

We provide an overview of the techniques in our proofs. We first mention a lemma below which should be intuitive.

Lemma 1. *Without loss of generality, we consider natural policies that use a maximal schedule to transmit at each time. Further, if a link is included in the schedule, its earliest-deadline packet will be selected for transmission.*

Proof. The proof is through an exchange argument. The details can be found in the technical report [20] \square

Frame Construction. A key step in the analysis of our scheduling algorithms is a careful frame construction. We emphasize that the frame construction is only for the purpose of analysis and is *not* part of our algorithms. The F-framed construction in [8] only works for i.i.d. arrivals and deadlines. Here, we need a construction that can handle our Markovian traffic model. We present this construction below where frames have random length as opposed to fixed length in [8].

Definition 4 (Frames and Cycles). *Starting from an initial traffic state $\tau(0) = \mathbf{x} \in \mathcal{X}$, let t_i denote the i -th return time of traffic Markov chain $\tau(t)$ to \mathbf{x} , $i = 1, \dots$. By convention, define $t_0 = 0$. The i -th cycle \mathcal{C}_i is defined from the beginning of time slot $t_{i-1}+1$ until the end of time slot t_i , with cycle length $C_i = t_i - t_{i-1}$. Given a fixed $k \in \mathbb{N}$, we define the i -th frame $\mathcal{F}_i^{(k)}$ as k consecutive cycles $\mathcal{C}_{(i-1)k+1}, \dots, \mathcal{C}_{ik}$, i.e., from the beginning of slot $t_{(i-1)k}+1$ until the end of slot t_{ik} . The length of the i -th frame is denoted by $F_i^{(k)} = \sum_{j=(i-1)k+1}^{ik} C_j$. Define $\mathcal{J}(\mathcal{F}^{(k)})$ to be the space of all possible traffic patterns $(\tau(t), t \in \mathcal{F}^{(k)})$ during a frame $\mathcal{F}^{(k)}$. Note that these patterns start after \mathbf{x} and end with \mathbf{x} .*

By the strong Markov property and the positive recurrence of traffic Markov chain, frame lengths $F_i^{(k)}$ are i.i.d with mean $\mathbb{E}[F^{(k)}] = k\mathbb{E}[C]$, where $\mathbb{E}[C]$ is the mean cycle length which is a bounded constant [21]. In fact, since state space \mathcal{X} is finite, all the moments of C (and $F^{(k)}$) are finite. We choose a fixed k , and, when the context is clear, drop the dependence on k in the notation.

Define the class of *non-causal* \mathcal{F} -framed policies $\mathcal{P}_{NC}(\mathcal{F})$ to be the policies that, at the beginning of each frame \mathcal{F}_i , have

complete information about the traffic pattern in that frame, but have a restriction that they drop the packets that are still in the buffer at the end of the frame. Note that the number of such packets is at most $d_{\max}a_{\max}K$, which is negligible compared to the average number of packets in the frame, $\bar{a}_l\mathbb{E}[F] = \bar{a}_l k\mathbb{E}[C]$, as $k \rightarrow \infty$. Define the rate region

$$\Lambda_{NC}(\mathcal{F}) = \bigcup_{\mu \in \mathcal{P}_{NC}(\mathcal{F})} \Lambda_\mu. \quad (14)$$

Given a policy $\mu \in \mathcal{P}_{NC}(\mathcal{F})$, the time-average service rate \bar{I}_l of link l is well defined. In fact, by the renewal reward theorem (e.g. [26], Theorem 5.10), and boundedness of $\mathbb{E}[F]$,

$$\lim_{t \rightarrow \infty} \frac{\sum_{s=1}^t I_l(s)}{t} = \frac{\mathbb{E}[\sum_{t \in \mathcal{F}} I_l(t)]}{\mathbb{E}[F]} = \bar{I}_l. \quad (15)$$

Similarly for the deficit arrival rate λ_l , defined in (5),

$$\frac{\mathbb{E}[\sum_{t \in \mathcal{F}} \tilde{a}_l(t)]}{\mathbb{E}[F]} = \lambda_l, \quad l \in \mathcal{K}. \quad (16)$$

In Definition 4, each frame consists of k cycles. Using similar arguments as in [8], it is easy to see (and it is intuitive) that

$$\liminf_{k \rightarrow \infty} \Lambda_{NC}(\mathcal{F}^{(k)}) \supseteq \text{int}(\Lambda).$$

Hence, if we prove that for a causal policy ALG, there exists a constant ρ , and a large k_0 , such that for all $k \geq k_0$,

$$\rho \text{int}(\Lambda_{NC}(\mathcal{F}^{(k)})) \subseteq \Lambda_{ALG}, \quad (17)$$

then it follows that $\Lambda_{ALG} \supseteq \rho \text{int}(\Lambda)$. For our algorithms, we find a ρ such that (17) holds for *any traffic process* under our model. Then it follows that $\gamma_{ALG}^* \geq \rho$.

We define the *gain* of a policy μ at time t as

$$\mathcal{G}_\mu(t) = \sum_{l \in \mathcal{K}} w_l^\mu(t) I_l^\mu(t), \quad (18)$$

and the gain over a frame is $\sum_{t \in \mathcal{F}} \mathcal{G}_\mu(t)$. To prove (17), we rely on comparing the gain (total deficit of packets transmitted) by ALG and an optimal max-gain non-causal policy over a frame. The following proposition states the result for any general interference graph.

Proposition 1. *Consider a frame $\mathcal{F} \equiv \mathcal{F}^{(k)}$, for some fixed k based on returns of traffic process $\tau(t)$ to a state \mathbf{x} . Let $\|w(t_0)\| = \sum_{l \in \mathcal{K}} w_l(t_0)$ be the norm of the initial deficit vector at the start of the frame. Suppose for a causal policy ALG, given any $\epsilon > 0$, there is a W' such that when $\|w(t_0)\| > W'$,*

$$\frac{\mathbb{E}[\sum_{t \in \mathcal{F}} \mathcal{G}_{ALG}(t) | \mathcal{S}(t_0)]}{\mathbb{E}[\sum_{t \in \mathcal{F}} \mathcal{G}_{\mu^*}(t) | \mathcal{S}(t_0)]} \geq \rho - \epsilon, \quad (19)$$

where $\mathcal{S}(t_0) = (\Psi(t_0), w(t_0), \tau(t_0))$, and μ^* is the optimal non-causal policy that maximizes the gain over the frame. Then for any $\lambda \in \rho \text{int}(\Lambda_{NC}(\mathcal{F}))$, the network state process $\{\mathcal{S}(t)\}$ is positive recurrent, and further, the deficit queues are bounded in the sense of (6).

Gain Analysis. With Proposition 1 in hand, we analyze the achievable gain of our algorithm over a frame, compared with that of the optimal non-causal policy μ^* . Since characterizing

μ^* is hard, we extend a coupling technique from [16]–[18], [27] (developed for constant-weight single buffer analysis) to stochastic process $(\Psi(t), w(t), \tau(t))$ in a general network.

Consider a state $(\Psi(t), w(t), \tau(t))$ under our randomized algorithms at time $t \in \mathcal{F}$, and the state $(\Psi^{\mu^*}(t), w^{\mu^*}(t), \tau(t))$ under the optimal policy μ^* . We change the state of μ^* (by modifying its buffers and deficits) to make it identical to $(\Psi(t), w(t), \tau(t))$, but also give μ^* a larger gain $\mathcal{G}'_{\mu^*}(t) \geq \mathcal{G}_{\mu^*}(t)$ that can ensure the change is advantageous for μ^* considering the rest of the frame. Then, taking the expectation $\mathbb{E}[\mathcal{G}'(t)]$ with respect to the random decisions of our algorithm, AMIX-ND or AMIX-MS, and traffic patterns in a frame, we can bound the optimal gain of μ^* . Then we can prove the main results in view of Proposition 1.

V. PROOFS OF MAIN RESULTS

We first provide the proof of Proposition 1 and then provide the gain analysis of our algorithms. In what follows, we define

$$w_{max}(t) = \max_{l \in \mathcal{K}} w_l(t) \mathbb{1}(\Psi_l \neq 0), \quad (20)$$

to be the maximum deficit of a nonempty link at time t . Also define $[N] := \{1, 2, \dots, N\}$. We use $\mathbb{E}_X[\cdot]$ to denote conditional expectation $\mathbb{E}[\cdot|X]$. $\mathbb{E}^Y[\cdot]$ is used to explicitly indicate that expectation is taken with respect to some random variable Y . $|A|$ is used to denote the cardinality of set A .

A. Proof of Proposition 1

We look at the state process $\{\mathcal{S}(t)\}$ at times t_i when frames start. We show that this sampled chain is positive recurrent and further its mean deficit size is stable in the sense of (6). From this it follows that the original process $\{\mathcal{S}(t)\}$ is also stable as the mean frame size $\mathbb{E}[F]$ is bounded and the mean deficits within a frame can change at most by $a_{max}K\mathbb{E}[F]$.

Since $\lambda \in \rho \text{int}(\Lambda_{NC})$, we have for some $\epsilon > 0$, and some policy $\mu \in \mathcal{P}_{NC}(\mathcal{F})$,

$$\lambda \mathbb{E}[F](1 + 2\epsilon) \preceq \rho \mathbb{E}[\sum_{t \in \mathcal{F}} I^\mu(t)], \quad (21)$$

where \preceq is the component-wise inequality between vectors. This is simply due to the fact that in each frame, the number of deficit arrivals $\sum_{t \in \mathcal{F}} \tilde{a}(t)$ and the number of departures under the policy μ are i.i.d across the frames, with means $\mathbb{E}[F]\lambda$ and $\mathbb{E}[\sum_{t \in \mathcal{F}} I^\mu(t)]$, respectively, by the renewal reward theorem. Hence, to ensure stability, (21) must hold. Next, consider the Lyapunov function $V(t) := V(\mathcal{S}(t)) = \frac{1}{2} \sum_{l \in \mathcal{K}} w_l^2(t)$. Let $\{I(t), t \in \mathcal{F}\}$ denote the scheduling decisions by ALG within the frame. Using (4), we get

$$\begin{aligned} V(t_0 + F) - V(t_0) &= \frac{1}{2} \sum_{l \in \mathcal{K}} (w_l^2(t_0 + F) - w_l^2(t_0)) \\ &= \frac{1}{2} \sum_{t \in \mathcal{F}} \sum_{l \in \mathcal{K}} (w_l^2(t+1) - w_l^2(t)) \\ &\leq K a_{max}^2 F / 2 + \sum_{t \in \mathcal{F}} \sum_{l \in \mathcal{K}} w_l(t) (\tilde{a}_l(t) - I_l(t)). \end{aligned} \quad (22)$$

Let $\mathbb{E}_{t_0}[\cdot] = \mathbb{E}[\cdot|\mathcal{S}(t_0)]$. Then, over a frame,

$$\begin{aligned} \mathbb{E}_{t_0} [V(t_0 + F) - V(t_0)] &\leq \\ \mathbb{E}_{t_0} [\sum_{t \in \mathcal{F}} \sum_{l \in \mathcal{K}} w_l(t) \tilde{a}_l(t)] - \mathbb{E}_{t_0} [\sum_{t \in \mathcal{F}} \sum_{l \in \mathcal{K}} w_l(t) I_l(t)] + C_1, \end{aligned} \quad (23)$$

where $C_1 = K a_{max}^2 \mathbb{E}[F]/2$. Noting that $w_l(t_0) - F \leq w_l(t) \leq w_l(t_0) + a_{max}F$, at any $t \in \mathcal{F}$, we can bound

$$\mathbb{E}_{t_0} [\sum_{t \in \mathcal{F}} \sum_{l \in \mathcal{K}} w_l(t) \tilde{a}_l(t)] \leq \sum_{l \in \mathcal{K}} (w_l(t_0) \lambda_l \mathbb{E}[F]) + C_2, \quad (24)$$

where we have used (16) in the last inequality, and $C_2 = a_{max}^2 \mathbb{E}[F^2]K < \infty$. Let $I^*(t)$ be the scheduling decisions by the policy μ^* , and $I^\mu(t)$ be the scheduling decisions by the policy $\mu \in \mathcal{P}_{NC}(\mathcal{F})$ in (21). Note that μ^* is the optimal non-causal policy that maximizes the gain over the frame and can transmit packets from a previous frame (included in the initial buffer $\Psi(t_0)$). This only improves the performance of μ^* , compared to starting with empty buffers, hence,

$$\mathbb{E}_{t_0} [\sum_{t \in \mathcal{F}} \sum_{l \in \mathcal{K}} w_l^*(t) I_l^*(t)] \geq \mathbb{E}_{t_0} [\sum_{l \in \mathcal{K}} \sum_{t \in \mathcal{F}} w_l^\mu(t) I_l^\mu(t)]. \quad (25)$$

Using (25) and the proposition assumption, given $\epsilon > 0$, there is a W' such that, if $\|w(t_0)\| > W'$,

$$\begin{aligned} \mathbb{E}_{t_0} [\sum_{t \in \mathcal{F}} \sum_{l \in \mathcal{K}} w_l(t) I_l(t)] &\geq (\rho - \epsilon) \mathbb{E}_{t_0} [\sum_{t \in \mathcal{F}} \sum_{l \in \mathcal{K}} w_l^*(t) I_l^*(t)] \\ &\geq (\rho - \epsilon) \mathbb{E}_{t_0} [\sum_{l \in \mathcal{K}} \sum_{t \in \mathcal{F}} w_l^\mu(t) I_l^\mu(t)] \\ &\geq (\rho - \epsilon) \mathbb{E}_{t_0} [\sum_{l \in \mathcal{K}} \sum_{t \in \mathcal{F}} (w_l(t_0) - F) I_l^\mu(t)] \\ &\geq (\rho - \epsilon) \mathbb{E}_{t_0} [\sum_{l \in \mathcal{K}} \sum_{t \in \mathcal{F}} w_l(t_0) I_l^\mu(t)] - C_3, \end{aligned} \quad (26)$$

where $C_3 = K \mathbb{E}[F^2]$ is a constant. Using (26), (24), (23),

$$\begin{aligned} \mathbb{E}_{t_0} [(t_0 + F) - V(t_0)] &\leq C_4 + \sum_{l \in \mathcal{K}} \mathbb{E}[F] w_l(t_0) \lambda_l - (\rho - \epsilon) \sum_{l \in \mathcal{K}} w_l(t_0) \mathbb{E}_{t_0} [\sum_{t \in \mathcal{F}} I_l^\mu(t)] \\ &\leq C_4 + \sum_{l \in \mathcal{K}} w_l(t_0) \left(\lambda_l \mathbb{E}[F] - (\rho - \epsilon) \mathbb{E}_{t_0} [\sum_{t \in \mathcal{F}} I_l^\mu(t)] \right) \\ &\leq C_4 - \epsilon \mathbb{E}[F] \sum_{l \in \mathcal{K}} \lambda_l w_l(t_0), \end{aligned} \quad (27)$$

where $C_4 = C_1 + C_2 + C_3$, and in the last inequality we have used (21). Hence, given any $\delta > 0$, $\mathbb{E}_{t_0} [V(t_0 + F) - V(t_0)] \leq -\delta$ if

$$\|w(t_0)\| \geq \max((C_4 + \delta)/(\epsilon \mathbb{E}[F] \lambda_{min}), W'),$$

where $\lambda_{min} = \min_l \lambda_l$. This proves that the network Markov chain is positive recurrent by the Foster-Lyapunov Theorem and further the stability in the mean sense (6) follows [28] (note that the component $\Psi(t)$ lives in a finite state space).

B. Gain Analysis of AMIX-ND in Collocated Networks

Consider a subclass \mathcal{P}_{ND} of all the policies that schedule Non-Dominated (ND) links at each slot (recall Definition 1). We refer to policies in \mathcal{P}_{ND} as *ND-policies*. We show that the optimal ND-policy is close to the optimal non-restricted policy as stated below.

Lemma 2. *Consider any policy μ for scheduling packets in a frame \mathcal{F} . Then there is an ND-policy $\hat{\mu} \in \mathcal{P}_{ND}$ such that, under the same pattern $J \in \mathcal{J}(\mathcal{F})$ and initial state $\mathcal{S}(t_0)$,*

$$\sum_{t \in \mathcal{F}} \mathcal{G}_{\hat{\mu}}(t) \geq \sum_{t \in \mathcal{F}} \mathcal{G}_{\mu}(t) - a_{max}F^2$$

where F is the length of the frame.

Proof. Suppose the first time μ does not schedule a non-dominated link is t_0 . Suppose μ sends earliest-deadline packet $(w_y(t_0), d_y)$ from link y and $(w_x(t_0), d_x)$ be the earliest-deadline packet at a link x ($x \neq y$) that strictly dominates y , i.e. $w_x(t_0) \geq w_y(t_0)$, $d_x \leq d_y$. Consider some alternative policy μ' which has the same transmissions as μ up to time t_0 but transmits the packet of x at time t_0 instead. Let $w'_l(t)$, $l \in \mathcal{K}$ denote the link deficits under μ' . Note that $w'_l(t) = w_l(t)$, $\forall t \leq t_0$. We differentiate between 2 cases:

1) μ does not transmit packet x in the remaining time slots.

In this case, let μ' transmit the same packets as μ in the remaining slots (after t_0). We can show that the difference of the collected gain, as a result of the different deficits of the users can be bounded below

$$\sum_{t \in \mathcal{F}} \mathcal{G}_{\mu'}(t) - \sum_{t \in \mathcal{F}} \mathcal{G}_{\mu}(t) \geq -F$$

2) μ transmits packet x at some time slot t_a where $t_0 < t_a < t_0 + d_x$. In this case we let μ' transmit the same packets as μ for all $t > t_0$ except for time slot t_a in which it transmits packet y instead, which has not expired by the domination inequality $d_y \geq d_x$. In this case, we can show that

$$\sum_{t \in \mathcal{F}} \mathcal{G}_{\mu'}(t) - \sum_{t \in \mathcal{F}} \mathcal{G}_{\mu}(t) \geq -a_{max}F. \quad (28)$$

The detailed proof can be found in the technical report [20]. By repeating this process (at most F times), we can transform μ to $\hat{\mu}$. From this, the final result follows. \square

Lemma 3. *For each slot $t \in \mathcal{F}$, the gain obtained by AMIX-ND, and the amortized gain by any ND-policy $\hat{\mu}$, starting from some state $\mathcal{S}(t)$ satisfy:*

$$\mathbb{E}^R[\mathcal{G}'_{\hat{\mu}}(t) | \mathcal{S}(t)] \leq w_{max}(t) + \mathcal{E}_0 \quad (29)$$

$$\mathbb{E}^R[\mathcal{G}_{\text{AMIX-ND}}(t) | \mathcal{S}(t)] \geq w_{max}(t)\rho \quad (30)$$

where $\rho = \left(1 - \left(1 - \frac{1}{K}\right)^K\right)$ and $\mathcal{E}_0 = a_{max}d_{max} + 2F$, and $\mathbb{E}^R[\cdot]$ is expectation with respect to the random decisions of AMIX-ND.

Proof. At time t , after the new arrivals have happened, we have state $\mathcal{S}(t)$. AMIX-ND decides probabilistically to transmit a packet (w_f, e_f) from a non-dominated link $f \in \mathcal{B}_{ND}(t)$, and the ND-policy $\hat{\mu}$ transmits a packet (w_z, e_z) from some

other link z . We distinguish two cases following the same method as in [18] but for time-varying weights.

1) $e_f \leq e_z, w_f \leq w_z$: To maintain the same buffers for both algorithms, we remove the packet e_f from the buffer of link f under $\hat{\mu}$ and inject the packet with deadline e_z to link z so that $\hat{\mu}$ gets a packet with higher deadline and higher weight at the time t . Since both packets will expire in at most d_{max} slots, the deficit of f can only increase by at most $d_{max}a_{max}$ before packet e_f expires. Therefore giving $\hat{\mu}$ this additional compensation will guarantee that the modification is advantageous. Further, we decrease the deficit from link f by one ($w_f - 1$ in $\hat{\mu}$) and we increase the deficit of link z by one ($w_z + 1$ in $\hat{\mu}$). Then $\hat{\mu}$ and AMIX-ND have the same exact state. Making this change in the deficit will reduce the gain for each packet transmitted from link f in the future by one. To compensate for this, we give $\hat{\mu}$ extra gain which is the number of packets transmitted from link f for the rest of the frame, which is less than F . Hence, the total compensation is bounded by $F + a_{max}d_{max}$.

2) $e_z \leq e_f, w_z \leq w_f$: In this case, we allow $\hat{\mu}$ to additionally transmit the packet e_f at time t , and inject a copy of packet e_z to the buffer of link z . This makes the buffers identical, but results in the decrease of deficit of link f by one, which might not be advantageous for $\hat{\mu}$ for future times. To guarantee that the change is advantageous for $\hat{\mu}$, we give it one extra reward for each possible transmission from link f in the rest of the frame, which is less than F .

Let $\mathcal{G}'_{\hat{\mu}}^{(h_i)}(t)$ denote the reward (including the compensation) gained by $\hat{\mu}$ when it transmits a non-dominated packet h_i (recall h_i from Algorithm 1). Then

$$\begin{aligned} \mathbb{E}^R[\mathcal{G}'_{\hat{\mu}}^{(h_i)}(t) | \mathcal{S}^t] &= \sum_{h_j: j < i} p_{h_j}(t) (w_{h_j}(t) + F) \\ &\quad + w_{h_i}(t) + F + a_{max}d_{max} \\ &\leq w_{h_i}(t) + \sum_{h_j: j < i} p_{h_j}(t) w_{h_j}(t) + \mathcal{E}_0 \end{aligned} \quad (31)$$

where $\mathcal{E}_0 = a_{max}d_{max} + 2F$. Using the assigned probabilities (line 4 in Algorithm 2), it is easy to verify that (31) attains its maximum for $i = 1$, which is equal to $w_{h_1}(t) + \mathcal{E}_0 = w_{max}(t) + \mathcal{E}_0$. Hence, (29) indeed holds.

Now regarding AMIX-ND, similar derivation applies as in [19] to get the final bound. To see that, first let the number of links with positive probability be $B \leq K$. Then

$$\begin{aligned} \mathbb{E}^R[\mathcal{G}_{\text{AMIX-ND}}(t) | \mathcal{S}^t] &= \sum_{i \in [B]} w_{h_i}(t) p_{h_i}(t) = \\ &\sum_{i \in [B-1]} w_{h_i}(t) p_{h_i}(t) + \left(1 - \sum_{i \in [B-1]} p_{h_i}(t)\right) w_{h_B}(t) \stackrel{(a)}{=} \\ &w_{h_1}(t) \left(1 - \prod_{i=1}^{B-1} (1 - p_{h_i}(t)) \sum_{i=1}^{B-1} p_{h_i}(t)\right) \stackrel{(b)}{\geq} \\ &w_{h_1}(t) \left(1 - \left(\frac{B-1}{B}\right)^B\right), \end{aligned}$$

where (a) follows from the form of probabilities, and (b) follows by applying the inequality between arithmetic and geometric means of B terms: $(1 - p_{h_i}(t))$, $i \in [B - 1]$, and $\sum_{i=1}^{B-1} p_{h_i}(t)$. \square

Lemma 4. *Over any frame \mathcal{F} , with initial state $\mathcal{S}(t_0) = (\Psi(t_0), w(t_0), \tau(t_0))$, and any ND-policy $\hat{\mu}$.*

$$\lim_{\|w(t_0)\| \rightarrow \infty} \frac{\mathbb{E}^{R,J}[\sum_{t \in \mathcal{F}} \mathcal{G}_{\text{AMIX-ND}}(t) | \mathcal{S}(t_0)]}{\mathbb{E}^J[\sum_{t \in \mathcal{F}} \mathcal{G}_{\hat{\mu}}(t) | \mathcal{S}(t_0)]} \geq \rho. \quad (32)$$

Proof. Given the initial state $\mathcal{S}(t_0)$ and frame size F , consider all the traffic patterns of length F . Taking expectations of the result of Lemma 3, with respect to random traffic patterns J of length F and the decisions of ALG, we get

$$\begin{aligned} \mathbb{E}^{R,J}[\mathbb{E}^R[\mathcal{G}'_{\hat{\mu}}(t) | \mathcal{S}(t)] | \mathcal{S}(t_0), F] &\leq \mathbb{E}^{R,J}[w_{\max}(t) | \mathcal{S}(t_0), F] + \mathcal{E}_0 \\ \mathbb{E}^{R,J}[\mathbb{E}^R[\mathcal{G}_{\text{ALG}}(t) | \mathcal{S}(t)] | \mathcal{S}(t_0), F] &\geq \mathbb{E}^{R,J}[w_{\max}(t) | \mathcal{S}(t_0), F] \rho \end{aligned}$$

where $\text{ALG} = \text{AMIX-ND}$. Now notice that

$$\begin{aligned} \mathbb{E}^{R,J}[\mathbb{E}^R[\mathcal{G}'_{\hat{\mu}}(t) | \mathcal{S}(t)] | \mathcal{S}(t_0), F] &= \\ \mathbb{E}^{R,J}[\mathbb{E}^R[\mathcal{G}'_{\hat{\mu}}(t) | \mathcal{S}(t), \mathcal{S}(t_0)] | \mathcal{S}(t_0), F] &= \mathbb{E}^{R,J}[\mathcal{G}'_{\hat{\mu}}(t) | \mathcal{S}(t_0), F] \end{aligned}$$

where the first equality is due to the fact that, given $\mathcal{S}(t)$ and F , the gain of $\hat{\mu}$ at time t depends on current $\mathcal{S}(t)$ and future traffic pattern in the frame, but not on the past. The second equality is by the tower property of conditional expectation. Therefore, we get

$$\mathbb{E}^{R,J}[\mathcal{G}'_{\hat{\mu}}(t) | \mathcal{S}(t_0), F] \leq \mathbb{E}^{R,J}[w_{\max}(t) | \mathcal{S}(t_0), F] + \mathcal{E}_0 \quad (33)$$

Using similar arguments for the expected gain of AMIX-ND,

$$\mathbb{E}^{R,J}[\mathcal{G}_{\text{ALG}}(t) | \mathcal{S}(t_0), F] \geq \mathbb{E}^{R,J}[w_{\max}(t) | \mathcal{S}(t_0), F] \rho \quad (34)$$

The result then follows by summing the gains over time slots in the frame, taking the expectation with respect to frame size F , using the inequality between amortized gain and actual gain, and also showing that $\lim_{\|w(t_0)\| \rightarrow \infty} \mathbb{E}^{R,J}[\sum_{t \in \mathcal{F}} w_{\max}(t) | \mathcal{S}(t_0)] = \infty$. The details can be found in the technical report [20]. \square

Theorem 3. *For any policy μ , and AMIX-ND, given any $\epsilon > 0$, there is W' such that when $\|w(t_0)\| \geq W'$:*

$$\mathbb{E}_{\mathcal{S}(t_0)}\left[\sum_{t \in \mathcal{F}} \mathcal{G}_{\text{AMIX-ND}}(t)\right] \geq (\rho - \epsilon) \mathbb{E}_{\mathcal{S}(t_0)}\left[\sum_{t \in \mathcal{F}} \mathcal{G}_{\mu}(t)\right]$$

Proof. The proof follows from Lemma 2 and Lemma 4 and noting that $\mathbb{E}_{t_0}[\sum_{t \in \mathcal{F}} \mathcal{G}_{\mu}(t)] \rightarrow \infty$, as $\|w(t_0)\| \rightarrow \infty$. The details can be found in the technical report [20]. \square

C. Gain Analysis of AMIX-MS in General Networks

First we show that binary search in Algorithm 3 suffices for computing \bar{n} defined in (13).

Proposition 2. *The binary search in Algorithm 3 computes \bar{n} as defined in (13).*

Proof. The proof can be found in the technical report [20]. \square

We next state Lemmas 5 and 6 regarding the properties of the probabilities used by AMIX-MS, which are used in the

gain analysis. The proofs follow directly from the probabilities used by AMIX-MS (see the technical report [20]).

Lemma 5. *$C_n(t)$ (defined in (11)) is strictly decreasing as a function of n , for $\bar{n} \leq n \leq |\mathcal{M}|$.*

Lemma 6. *If $i \notin [\bar{n}]$ and $j \in [\bar{n}]$, for the choice of probabilities $p_k^{\bar{n}}(t)$ in (12) selected by AMIX-MS, we have*

$$\begin{aligned} W_{M_i}(t) + \sum_{k \in [\bar{n}]} p_k^{\bar{n}}(t) W_{M_k}(t) &< \\ W_{M_j}(t) + \sum_{k \in [\bar{n}] \setminus \{j\}} p_k^{\bar{n}}(t) W_{M_k}(t) \end{aligned}$$

Lemma 7. *For each time $t \in \mathcal{F}$, the gain obtained by AMIX-MS, and the amortized gain obtained by the Max-Gain policy μ , starting from some state $\mathcal{S}(t)$, satisfy:*

$$\mathbb{E}^R[\mathcal{G}'_{\mu}(t) | \mathcal{S}(t)] \leq \sum_{i \in [\bar{n}]} W_{M_i}(t) - (\bar{n} - 1) C_{\bar{n}}(t) + \mathcal{E}_m \quad (35)$$

$$\mathbb{E}^R[\mathcal{G}_{\text{AMIX-MS}}(t) | \mathcal{S}(t)] = \sum_{i \in [\bar{n}]} W_{M_i}(t) - \bar{n} C_{\bar{n}}(t) \quad (36)$$

where $\mathcal{E}_m = KF$ and \mathbb{E}^R is with respect to decisions of AMIX-MS.

Proof. Using the probabilities computed by AMIX-MS, the expected gain of AMIX-MS at time t is

$$\mathbb{E}[\mathcal{G}_{\text{AMIX-MS}}(t)] = \sum_{i \in [\bar{n}]} p_i^{\bar{n}}(t) W_{M_i}(t) = \sum_{i \in [\bar{n}]} W_{M_i}(t) - \bar{n} C_{\bar{n}}(t)$$

Next, for the amortized gain of the Max-Gain Policy μ , we will apply the same technique as in the collocated networks case, where we modify the buffers and give μ additional reward. Suppose μ transmits M_i , and AMIX-MS transmits some M_j . We make the buffers the same by allowing μ to additionally transmit all the packets that are transmitted by AMIX-MS but not by μ (i.e., in links $M_j \setminus M_i$). Since this will result in a decrease of the deficit by one for each link in $M_j \setminus M_i$ for μ in the remaining slots, we give μ an additional reward $\mathcal{E}_m = KF$ which is an upper bound on the number of packets transmitted by μ from links $M_j \setminus M_i$ in the remaining slots. To compute the expected gain, we consider two cases:

Case 1. $i \in [\bar{n}]$. In this case, we can write

$$\begin{aligned} \mathbb{E}[\mathcal{G}'_{\mu}^{M_i}(t)] &= W_{M_i}(t) + \sum_{j \in [\bar{n}] \setminus \{i\}} p_j^{\bar{n}}(t) (W_{M_j \setminus M_i}(t) + \mathcal{E}_m) \\ &\leq W_{M_i}(t) + \sum_{j \in [\bar{n}] \setminus \{i\}} p_j^{\bar{n}}(t) (W_{M_j}(t) + \mathcal{E}_m) \end{aligned} \quad (37)$$

$$\begin{aligned} &\leq W_{M_i}(t)(1 - p_i^{\bar{n}}(t)) + \sum_{j \in [\bar{n}]} p_j^{\bar{n}}(t) W_{M_j}(t) + \mathcal{E}_m \\ &= C_{\bar{n}}(t) + \sum_{i \in [\bar{n}]} W_{M_i}(t) - \bar{n} C_{\bar{n}}(t) + \mathcal{E}_m. \end{aligned} \quad (38)$$

Case 2. $i \notin [\bar{n}]$. In this case, we have

$$\begin{aligned} \mathbb{E}[\mathcal{G}'_{\mu}^{M_i}(t)] &\leq W_{M_i}(t) + \sum_{k \in [\bar{n}]} p_k^{\bar{n}}(t) (W_{M_k}(t) + \mathcal{E}_m) \\ &\stackrel{a}{<} W_{M_j}(t) + \sum_{k \in [\bar{n}] \setminus \{j\}} p_k^{\bar{n}}(t) W_{M_k}(t) + \mathcal{E}_m \\ &= C_{\bar{n}}(t) + \sum_{i \in [\bar{n}]} W_{M_i}(t) - \bar{n} C_{\bar{n}}(t) + \mathcal{E}_m, \end{aligned}$$

where in (a) we applied Lemma 6 for i, j . Note that in both cases, the upper bound is the same and does not depend on the particular choice of M_i . \square

Lemma 8. For $C_{\bar{n}}(t)$ in (11), We have

$$\frac{\sum_{i \in [\bar{n}]} W_{M_i}(t) - \bar{n}C_{\bar{n}}(t)}{\sum_{i \in [\bar{n}]} W_{M_i}(t) - (\bar{n}-1)C_{\bar{n}}(t)} \geq \frac{|\mathcal{M}|}{2|\mathcal{M}|-1}.$$

Proof. The result follows by inequality between arithmetic and harmonic means,

$$\frac{1}{\bar{n}} \sum_{i \in [\bar{n}]} W_{M_i}(t) \geq \frac{\bar{n}}{\sum_{i \in [\bar{n}]} W_{M_i}(t)^{-1}},$$

and noting that $|\mathcal{M}| \geq \bar{n}$. The details can be found in the technical report [20]. \square

Theorem 4. Under AMIX-MS, given any $\epsilon > 0$ there is W' such that for all $\|w_0\| = \sum_{l \in \mathcal{K}} w_L(t_0) \geq W'$,

$$\mathbb{E}_{t_0}^{R,J} \left[\sum_{t \in \mathcal{F}} \mathcal{G}_{\text{AMIX-MS}}(t) \right] \geq (\rho - \epsilon) \mathbb{E}_{t_0}^J \left[\sum_{t \in \mathcal{F}} \mathcal{G}_\mu(t) \right],$$

where μ is any non-causal policy, and $\rho = \frac{|\mathcal{M}|}{2|\mathcal{M}|-1}$.

Proof. By using Lemma 7, and summing and taking expectation similar to the proof of Lemma 4, it follows that

$$\begin{aligned} \mathbb{E} \left[\sum_{t \in \mathcal{F}} \mathcal{G}_{\text{AMIX-MS}}(t) \mid \mathcal{S}(t_0) \right] &= \mathbb{E} \left[\sum_{t \in \mathcal{F}} x(t) \mid \mathcal{S}(t_0) \right] \\ \mathbb{E} \left[\sum_{t \in \mathcal{F}} \mathcal{G}_{\hat{\mu}}(t) \mid \mathcal{S}(t_0) \right] &\leq \bar{\mathcal{E}}_m + \mathbb{E} \left[\sum_{t \in \mathcal{F}} y(t) \mid \mathcal{S}(t_0) \right] \end{aligned}$$

where $\bar{\mathcal{E}}_m = K \mathbb{E}[F^2]$, and $x(t) = y(t) - C_{\bar{n}}(t)$, where

$$y(t) := \sum_{i \in \bar{n}(t)} W_{M_i}(t) - (\bar{n}(t) - 1)C_{\bar{n}}(t).$$


It is easy to verify that $y(t) \geq W_{M_1}(t) \geq w_{\max}(t)$. Letting $\|w_0\| \rightarrow \infty$, and using Lemma 8 yields the result. The details can be found in the technical report [20]. \square

VI. SIMULATION RESULTS


In our simulations, we consider two cases for the deficit admission (see the model section): one is based on coin tossing where each arrival on a link l is counted as deficit with probability p_l , and the other is deterministic, where each arrival increases the deficit by exactly p_l .

We compare the performance of our randomized algorithms, AMIX-ND and AMIX-MS with LDF. Recall that LDF chooses the longest-deficit link, then removes the interfering links with this link, and repeats the procedure. We further consider two versions of LDF: One is LDF that does a random tie breaking when presented with a deficit tie (LDF-RD), and the other version tries to schedule the non-dominated link and its earliest-deadline packet (LDF-ED) in such tie situations. In the plots, we compare the average deficit (over all links) as we vary the value of the delivery ratio.

Collocated Networks. We first consider two interfering links with deterministic deficit admission. The traffic is periodic and consists of alternating Pattern A and Pattern B of Figure 1, with the delivery ratios satisfying $p_2 = p_1 + 0.001$. Figure 3a shows the result. As we can see, AMIX-ND is able to achieve



(a) Deterministic deficit admission (b) Coin-tossing deficit admission
Fig. 3: Comparison between AMIX-ND and LDF policies in a two-link network.



(a) Coin-tossing deficit admission (b) Deterministic deficit admission
Fig. 4: Comparison between AMIX-MS and LDF policies in a lightly connected interference graph with 5 links.

roughly $p_1 = 0.996$, whereas both versions of LDF become unstable for $p_1 = 0.5 + \epsilon$. In Figure 3b, again for two users, we used a traffic that consists of Pattern C followed by Pattern B, repeatedly. This time we keep $p_1 = p_2$. AMIX-ND achieved near $p_1 = 1.0$, whereas the better version of LDF achieved roughly 0.75, resulting in a gap of around 0.25.

General Networks. We consider an interference graph involving 5 links, with interference edges $E_l = \{(l_1, l_2), (l_2, l_3), (l_2, l_4), (l_4, l_5)\}$. For links l_2 and l_5 , we have a periodic traffic with period $t = 5$, where in slot 1 there are 2 packets arriving with deadline 2 and 3 and in slot 4 a packet arrives with deadline 1, and for links l_1, l_3, l_5 , we have 1 packet arriving with deadline 1 at slot 1, and 1 packet arriving with deadline 2 at slot 4. See Figure 4.

As we see, simulation results indicate that there are many scenarios that result in significant gap between our algorithms and LDF variants. This gap is especially pronounced when deterministic deficit admission is used, which is preferable as it provides a short-term guarantee on the deficit of a user.

VII. CONCLUSION

In this paper, we studied real-time traffic scheduling in wireless networks under an interference-graph model. Our results indicated the power of randomization over the prior deterministic greedy algorithms for scheduling real-time packets. In particular, our proposed randomized algorithms significantly outperform the well-known LDF policy in terms of efficiency ratio. As a future work, we will investigate efficient and distributed implementation of AMIX-MS for general graphs.

REFERENCES

[1] C. Lu, A. Saifullah, B. Li, M. Sha, H. Gonzalez, D. Gunatilaka, C. Wu, L. Nie, and Y. Chen, "Real-time wireless sensor-actuator networks for industrial cyber-physical systems," *Proceedings of the IEEE*, vol. 104, no. 5, pp. 1013–1024, 2015.

[2] J. Song, S. Han, A. Mok, D. Chen, M. Lucas, M. Nixon, and W. Pratt, "Wirelesschart: Applying wireless technology in real-time industrial process control," in *2008 IEEE Real-Time and Embedded Technology and Applications Symposium*. IEEE, 2008, pp. 377–386.

[3] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami, "Internet of things (iot): A vision, architectural elements, and future directions," *Future generation computer systems*, vol. 29, no. 7, pp. 1645–1660, 2013.

[4] I. Hou, V. Borkar, and P. R. Kumar, "A theory of QoS for wireless," in *Proc. IEEE International Conference on Computer Communications (INFOCOM)*, Rio de Janeiro, Brazil, April 2009.

[5] I. Hou and P. R. Kumar, "Admission control and scheduling for QoS guarantees for variable-bit-rate applications on wireless channels," in *Proc. ACM international symposium on Mobile ad hoc networking and computing (MOBIHOC)*, New Orleans, Louisiana, May 2009.

[6] ———, "Scheduling heterogeneous real-time traffic over fading wireless channels," in *Proc. IEEE International Conference on Computer Communications (INFOCOM)*, San Diego, California, March 2010.

[7] J. Jaramillo and R. Srikant, "Optimal scheduling for fair resource allocation in ad hoc networks with elastic and inelastic traffic," in *Proc. IEEE International Conference on Computer Communications (INFOCOM)*, San Diego, California, March 2010.

[8] X. Kang, W. Wang, J. J. Jaramillo, and L. Ying, "On the performance of largest-deficit-first for scheduling real-time traffic in wireless networks," *IEEE/ACM Transactions on Networking*, vol. 24, no. 1, pp. 72–84, 2014.

[9] X. Kang, I.-H. Hou, L. Ying *et al.*, "On the capacity requirement of largest-deficit-first for scheduling real-time traffic in wireless networks," in *Proceedings of the 16th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. ACM, 2015, pp. 217–226.

[10] A. A. Reddy, S. Sanghavi, and S. Shakkottai, "On the effect of channel fading on greedy scheduling," in *2012 Proceedings IEEE INFOCOM*. IEEE, 2012, pp. 406–414.

[11] J. J. Jaramillo, R. Srikant, and L. Ying, "Scheduling for optimal rate allocation in ad hoc networks with heterogeneous delay constraints," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 5, pp. 979–987, 2011.

[12] C. Joo, X. Lin, and N. B. Shroff, "Understanding the capacity region of the greedy maximal scheduling algorithm in multihop wireless networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 17, no. 4, pp. 1132–1145, 2009.

[22] M. J. Neely, "Queue stability and probability 1 convergence via lyapunov optimization," *arXiv preprint arXiv:1008.3519*, 2010.

[13] A. Dimakis and J. Walrand, "Sufficient conditions for stability of longest-queue-first scheduling: Second-order properties using fluid limits," *Advances in Applied probability*, vol. 38, no. 2, pp. 505–521, 2006.

[14] B. Hajek, "On the competitiveness of on-line scheduling of unit-length packets with hard deadlines in slotted time," in *Proceedings of the 2001 Conference on Information Sciences and Systems*, 2001.

[15] A. Kesselman, Z. Lotker, Y. Mansour, B. Patt-Shamir, B. Schieber, and M. Sviridenko, "Buffer overflow management in qos switches," *SIAM Journal on Computing*, vol. 33, no. 3, pp. 563–583, 2004.

[16] F. Y. Chin, M. Chrobak, S. P. Fung, W. Jawor, J. Sgall, and T. Tichý, "Online competitive algorithms for maximizing weighted throughput of unit jobs," *Journal of Discrete Algorithms*, vol. 4, no. 2, pp. 255–276, 2006.

[17] M. Bienkowski, M. Chrobak, and Ł. Jeż, "Randomized competitive algorithms for online buffer management in the adaptive adversary model," *Theoretical Computer Science*, vol. 412, no. 39, pp. 5121–5131, 2011.

[18] Ł. Jeż, "One to rule them all: A general randomized algorithm for buffer management with bounded delay," in *European Symposium on Algorithms*. Springer, 2011, pp. 239–250.

[19] ———, "A universal randomized packet scheduling algorithm," *Algorithmica*, vol. 67, no. 4, pp. 498–515, 2013.

[20] C. Tsanikidis and J. Ghaderi, "On the Power of Randomization for Scheduling Real-Time Traffic in Wireless Networks," *arXiv e-prints*, p. arXiv:2001.05146, Jan 2020.

[21] E. B. Dynkin, *Theory of Markov processes*. Courier Corporation, 2012.

[23] J. Ghaderi and R. Srikant, "On the design of efficient CSMA algorithms for wireless networks," in *49th IEEE Conference on Decision and Control (CDC)*. IEEE, 2010, pp. 954–959.

[24] J. Ni, B. Tan, and R. Srikant, "Q-CSMA: Queue-length-based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks," *IEEE/ACM Transactions on Networking (ToN)*, vol. 20, no. 3, pp. 825–836, 2012.

[25] D. Shah and J. Shin, "Delay optimal queue-based CSMA," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 1. ACM, 2010, pp. 373–374.

[26] S. M. Ross, *Applied probability models with optimization applications*. Courier Corporation, 2013.

[27] Ł. Jeż, F. Li, J. Sethuraman, and C. Stein, "Online scheduling of packets with agreeable deadlines," *ACM Transactions on Algorithms (TALG)*, vol. 9, no. 1, p. 5, 2012.

[28] S. P. Meyn and R. L. Tweedie, "Stability of markovian processes i: Criteria for discrete-time chains," *Advances in Applied Probability*, vol. 24, no. 3, pp. 542–574, 1992.