Randomized Scheduling of Real-Time Traffic in Wireless Networks Over Fading Channels

Christos Tsanikidis, Javad Ghaderi
Electrical Engineering Department, Columbia University, New York, NY, USA

Abstract—**Despite the rich literature on scheduling algorithms for wireless networks, algorithms that can provide deadline guarantees on packet delivery for general traffic and interference models are very limited. In this paper, we study the problem of scheduling real-time traffic under a *conflict-graph interference model* with *unreliable links* due to channel fading. Packets that are not successfully delivered within their deadlines are of no value. We consider traffic (packet arrival and deadline) and fading (link reliability) processes that evolve as an *unknown* finite-state Markov chain. The performance metric is efficiency ratio which is the fraction of packets of each link which are delivered within their deadlines compared to that under the optimal (unknown) policy. We first show a conversion result that shows classical non-real-time scheduling algorithms can be ported to the real-time setting and yield a constant efficiency ratio, in particular, Max-Weight Scheduling (MWS) yields an efficiency ratio of $1/2$. We then propose randomized algorithms that achieve efficiency ratios strictly higher than $1/2$, by carefully randomizing over the maximal schedules. We further propose low-complexity and myopic distributed randomized algorithms, and characterize their efficiency ratio. Simulation results are presented that verify that randomized algorithms outperform classical algorithms such as MWS and GMS.**

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

I. Introduction

There has been vast research on scheduling algorithms in wireless networks which mostly focus on maximizing long-term throughput when packets have no strict delay constraints. Max-Weight Scheduling (MWS) policy has been shown to be throughput optimal in such settings, attaining any desired throughput vector in the feasible throughput region [1]. Further, greedy scheduling policies such as LQF [2], [3], or distributed policies such as CSMA [4]–[6] have been proposed that alleviate the computational complexity of MWS and achieve a certain fraction of the throughout region. However, in many emerging applications, such as Internet of Things (IoT), vehicular networks, and edge computing, delays and deadline guarantees on packet delivery play an important role [7]–[9], as packets that are not received within specific deadlines are of little or no value, and are typically discarded. This *discontinuity* in the packet value as a function of latency makes the problem significantly more challenging than traditional scheduling where packets do not have strict deadlines.

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

There is an increasing body of work attempting to address the above challenge, however they either assume a frame-based traffic model [10]–[14], relax the interference graph constraints [15], or use greedy scheduling approaches like LDF [16], [17]. In the frame-based traffic model, time is divided into frames, and packet arrivals and their deadlines during a frame are assumed to be *known* at the beginning of the frame, and deadlines are constrained by the frame's length [10]–[14]. Under such assumptions, the optimal solution in each frame is a Max-Weight schedule. Note that unless the traffic is restricted to be synchronized across the users, such solutions are non-causal. A more general cyclic traffic is considered in [18] for the collocated case, however the proposed solutions are either computationally prohibitive or they provide no performance guarantees. The optimal scheduling policy (and the real-time throughout region) for general traffic patterns and interference graphs is unknown and very difficult to characterize. Largest-Deficit-First (LDF) is a causal policy which extends the well-studied Largest-Queue-First (LQF) from traditional scheduling to real-time scheduling. The performance of LDF has been studied in terms of *efficiency ratio*, which is the fraction of the real-time throughput region guaranteed by LDF. Under *i.i.d.* packet arrivals and deadlines, with no fading, LDF was shown to achieve an efficiency ratio of at least $\frac{1}{\beta+1}$ [16], where β is the *interference degree* of the network (which is the maximum number of links that can be scheduled simultaneously out of a link and its neighboring links). Recently, the work [19] has shown that through randomization it is possible to design algorithms that can significantly improve the prior algorithms, in terms of both efficiency ratio and traffic assumptions. Specifically, [19] proposed two randomized scheduling policies: AMIX-ND for collocated networks with an efficiency ratio of at least $\frac{e-1}{e} \approx 0.63$, and AMIX-MS for general interference graphs with an efficiency ratio of at least $\frac{|\mathcal{I}|}{2|\mathcal{I}|-1} > \frac{1}{2}$ ($|\mathcal{I}|$ is the number of maximal independent sets of the graph). However, the complexity of AMIX-MS can be prohibitive for implementation in large networks. Moreover, intrinsic wireless channel fading has not been considered in [19], and packet transmission over a link is assumed to be always reliable.

In this work, we consider an interference graph model of wireless network subject to fading, where packet transmissions over links are unreliable. We consider a joint traffic (packet arrival and deadline) and fading (link's success probability) process that evolves as an *unknown* Markov chain over a

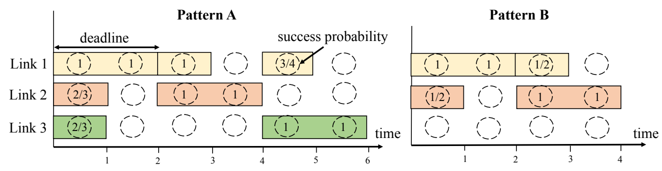

Fig. 1: An example of a Markovian traffic-and-fading process that alternates between two patterns. Each rectangle indicates a packet for a link. The left side of the rectangle corresponds to its arrival time, and its length corresponds to its deadline. The numbers in circles indicate the channel success probability of each link in each time slot.

finite state space. Can the existing traditional scheduling algorithms (which focus on long-term throughput with no deadline constraints) be used to provide guarantees for scheduling in this setting, *without making frame-based traffic assumptions*? We show that interestingly the answer is yes, but fading and deadlines might significantly degrade their performance.

Introducing fading in deadline-constrained scheduling makes the problem very complicated. For example, consider the initial two time slots in Pattern B in traffic-fading process of Figure 1 with two interfering links: link 1 has a packet with deadline 2 and link 2 has a packet with deadline 1. Link 1's success probability is 1 in the first two slots, and link 2's success probability is 0.5 in the first time slot. Not knowing the future traffic and fading, an opportunistic scheduler would prioritize link 1 in the first time slot and subsequently in the second time slot, but the optimal policy would always schedule link 2 in the first time slot and packet of link 1 in the second time slot. A key insight of our work is that careful randomization in decision is crucial to hedge against the risk of poor decision due to lack of the knowledge of future traffic and fading.

A. Contributions

Our main contributions can be summarized as follows.

- **Application of Traditional Scheduling Algorithms to Real-Time Scheduling.** Each non-empty link (i.e., with unexpired packets for transmission) is associated with a weight which is the product of its deficit counter and channel fading probability at that time, otherwise if the link is empty, the weight is zero. We show that any algorithm that provides a ψ-approximation to Max-Weight Schedule (MWS) under such weights, achieves an efficiency ratio of at least $\frac{\psi}{\psi+1}$ for real-time scheduling under any Markov traffic-fading process. As a consequence, MWS policy achieves an efficiency ratio of $\frac{1}{2}$, and GMS (Greedy Maximal Scheduling) provides an efficiency ratio of at least $\frac{1}{\beta+1}$.
- **Randomized Scheduling of Real-Time Traffic Over Fading Channels.** We extend [19] to show the power of randomization for scheduling real-traffic traffic over fading channels in general and collocated networks. By carefully randomizing over the maximal schedules, the algorithms can achieve an efficiency ratio of *at least* $\frac{|\mathcal{I}|}{2|\mathcal{I}|-1} > \frac{1}{2}$ in any general graph under any *unknown* Markov traffic-fading process. In the special case of a collocated network with i.i.d. channel success probability q, the algorithm can achieve an efficiency ratio of at least $(\frac{q}{1-e^{-q}}+1-q)^{-1}$, which ranges from 0.5 to 0.63.
- **Low-Complexity and Distributed Randomized Algorithms.** To address the high complexity of the randomized algorithm in general graphs, we propose a low-complexity covering-based randomization and a myopic distributed randomization. Given a coloring of the graph using χ colors, we can achieve an efficiency ratio of at least $\frac{1}{2\chi-1}$, by randomizing over χ schedules. Moreover, we show that a myopic distributed randomization, which is simple and easily implementable, can achieve an efficiency ratio of $\frac{1}{\Delta+1}$ in any graph with maximum degree Δ.

B. Notations

Some of the basic notations used in this paper are as follows. Given $\mathbf{y} \in \mathbb{R}^n$, $||\mathbf{y}|| = \sum_{i=1}^{n} |y_i|$. We use $\text{int}(A)$ to denote the interior of set A. $[x]^+ = \max\{x, 0\}$. $\mathbb{1}(E)$ is the indicator function of event E. $|A|$ is the cardinality of set A. $\mathbb{E}^Y[\cdot]$ is used to indicate that expectation is taken with respect to random variable Y.

II. Model and Definitions

Wireless Network and Interference Model. Consider a set of K links, denoted by the set $\mathcal{K}$. We assume time is divided into slots, and in every slot t, each link $l \in \mathcal{K}$ can attempt to transmit at most one packet. To model interference between links, we use the standard *interference graph* $G_I = (\mathcal{K}, E_I)$: Each vertex of G_I is a link, and there is an edge $(l_1, l_2) \in E_I$ if links l_1, l_2 interfere with each other. Hence, no two links that transmit packets can share an edge in G_I. Let $\mathcal{I}$ be the set of all maximal independent sets of graph G_I. Also let $\mathcal{B}(t) \subseteq \mathcal{K}$ denote the set of nonempty links (i.e., links that have packets available to transmit) at time t. We use $M(t)$ to denote the set of links scheduled at time t. By definition, $M(t)$ is a valid *schedule* if links in $M(t)$ are nonempty and form an independent set of G_I, i.e.,

$$M(t) \subseteq (\mathcal{B}(t) \cap D), \text{ for some } D \in \mathcal{I}. \tag{1}$$

A schedule is said to be maximal if no nonempty link can be added to the schedule without violating the interference constraints. In this case, '$\subseteq$' in (1) holds with '$=$'.

Fading Model. Transmission over a link is unreliable due to wireless channel fading. To capture the channel fading over link l, we use an ON-OFF model where link l at time t is ON ($C_l(t) = 1$) with probability $q_l(t)$, otherwise it is OFF ($C_l(t) = 0$). If scheduled, transmission over link l at time t is successful only if link l is ON. Let $\mathbf{q}(t) = (q_l(t), l \in \mathcal{K})$ be the vector of success probabilities of the links. We assume that at any time slot t, the link's success probability is known to the scheduler before making a decision. At the end of time slot t, link's transmitter receives a feedback from its receiver indicating whether transmission was successful or not. A special case of this model is when $q_l(t) \in \{0, 1\}$, $\forall l \in$

$\mathcal{K}$, in which case the channel state $C_l(t)$ is deterministically known to the scheduler, which requires periodic channel state estimation. Another special case is when $C_l(t)$ is i.i.d. with some probability q_l which eliminates the need for periodic estimation. Similar models have been used in [15], [20].

Overall, we define $I(t) = (I_l(t), l \in \mathcal{K})$ to denote the successful packet transmissions over the links at time t. Note that by definition, $I_l(t) = \mathbb{1}(l \in M(t))C_l(t)$, where $M(t)$ is the valid schedule at time t, and $C_l(t)$ is the channel state of link l.

Traffic Model. We assume a single-hop real-time traffic. We use $a_l(t) \leq a_{\max}$ to denote the number of packet arrivals at link l at time t, where $a_{\max} < \infty$ is a constant. Each arriving packet has a deadline which indicates the maximum delay that the packet can tolerate before successful transmission. A packet with deadline d at time t has to be successfully transmitted before the end of time slot $t + d - 1$, otherwise it will be discarded. We define the traffic process $\boldsymbol{\tau}(t) = (\tau_{l,d}(t), d = 1, \cdots, d_{\max}, l \in \mathcal{K})$, where $\tau_{l,d}(t)$ is the number of packets with deadline d arriving to link l at time t, and $d_{\max} < \infty$.

Traffic and Fading Process. In general, we assume that the joint traffic and fading process $\mathbf{z}(t) = (\mathbf{q}(t), \boldsymbol{\tau}(t))$ evolves as an "*unknown*" irreducible Markov chain over a finite state space $\mathcal{Z}$. See Figure 1 for an example of a Markovian traffic-fading process.

Without loss of generality, we make the following assumption to make this Markov chain *non-trivial*: For every link l, there are two states $\mathbf{z}^l, \mathbf{z'}^l \in \mathcal{Z}$ such that $\mathbf{z}^l$ has a packet arrival with deadline d, and $\mathbf{z'}^l$ has $q_l > 0$, and there is a positive probability that $\mathbf{z}(t)$ can go from $\mathbf{z}^l$ to $\mathbf{z'}^l$ in at most d time slots. This assumption simply states that it is possible to successfully transmit some packets of every link l within their deadlines. If a link does not satisfy this condition, we can simply remove it from the system. Note that $\mathbf{z}(t)$ is an irreducible finite-state Markov chain, hence it is positive recurrent [21], and time-average of any bounded function of $\mathbf{z}(t)$ is well defined, in particular the packet arrival rate for link l:

$$\lim_{t\to\infty} \tfrac{1}{t}\textstyle\sum_{s=1}^{t} a_l(s) =: \overline{a}_l. \tag{2}$$

Buffer Dynamics. The buffer of link l at time t, denoted by $\Psi_l(t)$, contains the existing packets at link l which have not expired yet and also the newly arrived packets at time t. The remaining deadline of each packet in $\Psi_l(t)$ 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 $\Psi_l(t)$. We also define $\Psi(t) = (\Psi_l(t); l \in \mathcal{K})$.

Delivery Requirement and Deficit. As in [10]–[13], [16], we assume that there is a minimum delivery ratio requirement p_l (QoS requirement) for each link $l \in \mathcal{K}$. This means we must successfully deliver at least p_l fraction of the incoming packets on each link l *within their deadlines*. Formally,

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

We define a deficit $w_l(t)$ which measures the number of successful packet transmissions owed to link l up to time t to fulfill its minimum delivery ratio. As in [13], [16], [19], the deficit evolves as

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

where $\widetilde{a}_l(t)$ 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 [13], [16], i.e., each packet arrival at link l increases the deficit by one with the probability p_l, and zero otherwise. We refer to $\widetilde{a}_l(t)$ as the *deficit arrival* process for link l. Note that it holds that

$$\lim_{t\to\infty} \tfrac{1}{t}\textstyle\sum_{s=1}^{t} \widetilde{a}_l(s) = \overline{a}_l p_l := \lambda_l, \ l \in \mathcal{K}. \tag{5}$$

We refer to λ_l as the deficit arrival rate of link l. Note that an 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 transmitted successfully from link l, i.e., $I_l(t) = 1$, the deficit is reduced by one. The dynamics in (4) define a deficit queueing system, with bounded increments/decrements, whose stability, e.g., in the sense

$$\limsup_{t\to\infty} \tfrac{1}{t}\textstyle\sum_{s=1}^{t} \mathbb{E}[w_l(s)] < \infty, \tag{6}$$

implies (3) holds [22]. Define the vector of deficits as $\mathbf{w}(t) = (w_l(t), l \in \mathcal{K})$. The system state at time t is then defined as

$$\mathcal{S}(t) = (\mathbf{q}(t), \boldsymbol{\tau}(t), \Psi(t), \mathbf{w}(t)). \tag{7}$$

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, deadlines, and channel success probabilities in order to make scheduling decisions. For a given traffic-fading process $\mathbf{z}(t)$, with fixed $\overline{a}_l$, defined in (2), 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 $\boldsymbol{\lambda} = (\lambda_l := \overline{a}_l p_l, l \in \mathcal{K})$ possible. For a given traffic process, we say the rate vector $\boldsymbol{\lambda} = (\lambda_l, l \in \mathcal{K})$ is supportable under some policy $\mu \in \mathcal{P}_C$ if all the deficit queues remain stable for that policy. Then one can define the supportable real-time rate region of the policy μ as

$$\Lambda_\mu = \{\boldsymbol{\lambda} \geq 0 : \boldsymbol{\lambda} \text{ is supportable by } \mu\}. \tag{8}$$

The supportable *real-time 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* defined as

$$\gamma^\star_\mu = \sup\{\gamma : \gamma\Lambda \subseteq \Lambda_\mu\}. \tag{9}$$

For a casual policy μ, we aim to provide a *universal lower bound* on the efficiency ratio that holds for "all" Markovian traffic-fading processes, *without* knowing the Markov chain.

For clarity, we will use $\mathbb{E}^{\mathrm{F}}[\cdot]$ to denote expectation with respect to the *random outcomes of the fading channel*, and $\mathbb{E}^{\mathrm{R}}[\cdot]$ to denote expectation with respect to the *random decisions of a randomized algorithm* ALG, whenever applicable.

III. Scheduling Algorithms and Main Results

Recall that $\mathcal{I}$ is the set of all maximal independent sets of interference graph G_I, and $\mathcal{B}(t)$ is the set of links that have packets to transmit at time t. At any time t, we define the set of maximal schedules $\mathcal{M}(t) = \{D \cap \mathcal{B}(t), D \in \mathcal{I}\}$.

Define the *gain* of maximal schedule $M \in \mathcal{M}(t)$ to be the total deficit of packets transmitted successfully at time t, i.e., $\mathcal{G}_M(t) := \sum_{l \in M} C_l(t) w_l(t)$. We define the *weight* of maximal schedule M to be its expected gain, conditional on $\mathcal{S}(t)$, i.e.,

$$W_M(t) := \mathbb{E}^{\mathrm{F}}\left[\mathcal{G}_M(t) | \mathcal{S}(t)\right] = \sum_{l \in M} w_l(t) q_l(t). \quad (10)$$

We use $M^\star(t) := \arg\max_{M \in \mathcal{M}(t)} W_M(t)$ to denote the *Max-Weight Schedule* (MWS) at time t, and $W^\star(t) := W_{M^\star(t)}(t)$ to denote its weight.

Given a Markov policy ALG, let $\pi_M(t)$ be the probability that ALG selects maximal schedule M at time t. Hence, the probability that link l is scheduled at time t is

$$\phi_l(t) = \textstyle\sum_{M \in \mathcal{M}(t)} \pi_M(t) \mathbb{1}(l \in M). \quad (11)$$

Further, the expected gain of ALG is given by

$$\begin{aligned} \mathbb{E}^{\mathrm{R,F}}\left[\mathcal{G}_{ALG}(t) | \mathcal{S}(t)\right] &= \sum_{M \in \mathcal{M}(t)} \pi_M(t) W_M(t) \quad (12) \\ &= \sum_{l \in \mathcal{K}} \phi_l(t) q_l(t) w_l(t). \quad (13) \end{aligned}$$

Without loss of generality, we consider natural policies that transmit the earliest-deadline packet from every selected link in the schedule. This is because, similar to [19], if a policy transmits a packet that is not the earliest-deadline, that packet can be replaced with the earliest-deadline packet of that link, and this only improves the state. Similarly, the optimal policy always selects a maximal schedule at any time.

A. Converting Classical Non-Real-Time Algorithms for Real-Time Scheduling

The following theorem allows us to convert non-real-time scheduling policies to real-time scheduling policies, hence enabling the use of numerous policies from the literature of traditional non-real-time scheduling. More specifically, policies whose expected gain at any time t is ψ-fraction of the gain of MWS (i.e. $W^\star(t)$) yield an efficiency ratio of at least $\psi/(\psi+1)$. The result is stated formally in Theorem 1,

Theorem 1. *Consider any policy* ALG *such that, at every time t, $\mathbb{E}^{\mathrm{R,F}}\left[\mathcal{G}_{ALG}(t) | \mathcal{S}(t)\right] \geq \psi W_{M^\star}(t)$, whenever $\|\mathbf{w}(t)\| \geq W'$, for some finite W'. Then*

$$\gamma^\star_\mu \geq \frac{\psi}{\psi + 1}.$$

Note that this conversion results in deadline-oblivious policies which can be preferred in cases where information about the deadlines of packets is either not accurate or not available. We remark that for certain policies the result of Theorem 1 is tight as seen in Corollary 1.1 below, whereas for other policies the bound can be loose.

Corollary 1.1. MWS *policy provides efficiency ratio $\gamma^\star_{\mathrm{MWS}} = \frac{1}{2}$ for real-time scheduling under Markov traffic-fading processes.*

Proof. Using Theorem 1 for MWS with $\psi = 1$, we directly obtain $\gamma^\star_{\mathrm{MWS}} \geq 1/2$. We can get the the opposite inequality through an adversarial example. If we consider a simple network with two interfering links without fading, then MWS reduces to LDF, which has been shown to have $\gamma^\star \leq 1/2$ for Markovian traffic with deterministic deficit admission [16], [19] (recall that efficiency ratio $\gamma^\star$ is defined as a universal bound for all traffic-fading processes for any given graph). □

Consider a Greedy Maximal Scheduling (GMS) policy defined as follows: Order the nonempty links $\mathcal{B}(t)$ in the decreasing order of the product $w_l(t)q_l(t)$. Then construct a schedule recursively by including the nonempty link with the largest $w_l(t)q_l(t)$, removing the interfering links, and repeating the same procedure over the remaining links.

The following corollary extends the result of [16] which was shown for i.i.d. traffic without fading for LDF.

Corollary 1.2. GMS *provides an efficiency ratio $\gamma^\star_{\mathrm{GMS}} \geq \frac{1}{\beta+1}$ for real-time scheduling under any Markov traffic-fading process, where β is the interference degree of the graph.*

Proof. It is straightforward to show that GMS provides a $1/\beta$ approximation to MWS, through standard arguments. The result then follows by applying Theorem 1 with $\psi = 1/\beta$. Detailed derivation can be found in the technical report [23]. □

Note that many other results from traditional scheduling literature can be converted using Theorem 1. For example, the DistGreedy policy [24] would also obtain $\gamma^\star_{\mathrm{DistGreedy}} \geq \frac{1}{\beta+1}$ but has a lower complexity than GMS.

B. Randomized Scheduling Algorithms

In this section, we extend AMIX policies for collocated networks and general networks introduced in [19] to incorporate fading. We refer to the generalized policies as FAMIX (*Fading-based Adaptive MIX*). We describe these generalized policies below.

1) FAMIX-MS: Randomized Scheduling in General Graphs: Let $R := |\mathcal{M}(t)|$ be the number of maximal schedules at time t. We index and order the maximal schedules such that $M^{(i)} \in \mathcal{M}(t)$ has the i-th largest weight (based on (10)), i.e.,

$$W_{M^{(1)}}(t) \geq W_{M^{(2)}}(t) \cdots \geq W_{M^{(R)}}(t). \quad (14)$$

Define the *subharmonic* average of weight of the first n maximal schedules, $n \leq R$, to be

$$C_n(t) = \frac{n-1}{\sum_{i=1}^{n} \left(W_{M^{(i)}}(t)\right)^{-1}}. \quad (15)$$

Then select schedule $M^{(i)}$ with probability

$$\pi^{\bar{n}}_{M^{(i)}}(t) \equiv \pi^{\bar{n}}_i(t) = \begin{cases} 1 - \dfrac{C_{\bar{n}}(t)}{W_{M^{(i)}}(t)}, & 1 \leq i \leq \bar{n} \\ 0, & \bar{n} < i \leq R \end{cases} \quad (16)$$

where $\bar{n}$ is the largest n such that $\{\pi_i^n(t), 1 \le i \le R\}$ defines a valid probability distribution, i.e., $\pi_i^n(t) \ge 0$, and $\sum_{i=1}^R \pi_i^n(t) = 1$.

Remark 1. It has been shown in [19] that $\bar{n}$ in such a distribution can be found through a binary search. Essentially the difference between FAMIX-MS and AMIX-MS in [19] lies in different definitions for the weight of a schedule.

Theorem 2. *In a general interference graph G_I with maximal independent sets $\mathcal{I}$, the efficiency ratio of* FAMIX-MS *is*

$$\gamma^\star_{\text{FAMIX-MS}} \ge \frac{|\mathcal{I}|}{2|\mathcal{I}|-1} > \frac{1}{2}.$$

Remark 2. Theorem 2 shows that with randomization we can do *strictly* better than MWS (Corollary 1.1). For example, in the case of a complete bipartite interference graph (where there are two maximal independent sets), FAMIX-MS yields an efficiency ratio of at least $2/3$, while MWS yields $1/2$.

2) FAMIX-ND: Randomized Scheduling in Collocated Graphs: Extending AMIX-ND [19] to fading channels is more challenging. In particular, the derivation in [19] relied on two main ideas: (1) it is sufficient to consider only a restricted set of "non-dominated" links for transmission, and (2) there is an ordering among the non-dominated links such that given two non-dominated links, having packet in buffer from one of the links is always preferred to that from the other link. Finding such domination relationship for a general Markov fading process is difficult. Here, we describe an extension under a simplified fading process, where $q_l(t) = q_l = q$, i.e., the links' success probabilities are fixed and equal. We allow the channel state across links to be either independent or positively correlated, i.e., $\Pr[C_{l_2}(t)=1 | C_{l_1}(t)=1] \ge q_{l_2}$. The above setting could be a reasonable approximation in collocated networks where links have similar reliabilities, and an active channel for one link implies a better condition for the overall shared wireless medium.

Let $e_l(t)$ denote the deadline of the earliest-deadline packet of link l at time t. We say that link l_1 dominates link l_2 at time t if $e_{l_1}(t) \le e_{l_2}(t)$, $w_{l_1}(t) \ge w_{l_2}(t)$, i.e., link l_1 is more urgent and has a higher deficit. Based on this definition, the set of non-dominated links at any time can be found though a simple recursive procedure as in [19], i.e., by adding the largest-deficit nonempty link, removing all the links dominated by it from consideration and repeating the process for the remaining links. The following theorem describes FAMIX-ND and its efficiency ratio for a collocated network with a channel success probability q.

Theorem 3. *Consider a collocated network, where $q_l = q$, $\forall l \in \mathcal{K}$, and channels of links are independent or positively correlated. Order and re-index the non-dominated links such that*

$$w_1(t) \ge w_2(t) \ge \dots \ge w_n(t).$$

Starting from $i = 1$, assign probability $\pi_i(t)$ to the i-th non-dominated link,

$$\pi_i(t) = \min\Big\{\frac{1}{q}\Big(1 - \frac{w_{i+1}(t)}{w_i(t)}\Big), 1 - \sum_{j<i} \pi_j(t)\Big\}. \tag{17}$$

FAMIX-ND *selects the i-th non-dominated link with probability $\pi_i(t)$ and transmits its earliest-deadline packet. Then*

$$\gamma^\star_{\text{FAMIX-ND}} \ge \Big(\frac{q}{1-e^{-q}} + (1-q)\Big)^{-1} := h(q). \tag{18}$$

Remark 3. Note that due to channel uncertainty, FAMIX-ND boosts the probability of larger-weight links. Intuitively, as q becomes small, deadlines of packets are "effectively" reduced, as each packet will need to be transmitted several times before success.

Remark 4. The lower bound $h(q)$ on efficiency ratio in (18) is a monotone function of q, which increases from 0.5 and to $\frac{e-1}{e}$, as q goes from 0 to 1. For $q = 1$, this recovers the result of [19] for non-fading channels, i.e., $\gamma^\star_{\text{AMIX-ND}} \ge \frac{e-1}{e}$.

In the case of unequal $q_i \in (q_{\min}, q_{\max})$, using (17) by replacing q with $q_{\min}$, will give an efficiency ratio of at least $(\frac{q_{\max}}{1-e^{-q_{\min}}} + 1 - q_{\max})^{-1} := h(q_{\min}, q_{\max})$. Depending on $q_{\max}, q_{\min}, K$, we can choose either FAMIX-ND or FAMIX-MS and achieve $\gamma^\star \ge \max\{h(q_{\min}, q_{\max}), \frac{K}{2K-1}\}$.

C. Low-Complexity and Distributed Randomized Variants

The general algorithm FAMIX-MS in Section III-B potentially randomizes over all the maximal schedules. This can be computationally expensive in large networks that may have many maximal schedules. In this section, we design variants that only need to consider a subset of the maximal schedules, or are distributed.

1) Covering-Based Randomized Algorithms: We propose variants that only need to consider a subset of the maximal schedules. Proposition 1 below states a sufficient condition under which randomization over a subset of the maximal schedules can provide a related approximation on the efficiency ratio.

Proposition 1. *Consider policy* ALG = FAMIX-MS$|_{\mathcal{M}_0(t)}$ *which, at any time t, selects a schedule from a subset $\mathcal{M}_0(t) \subseteq \mathcal{M}(t)$, according to probabilities of* FAMIX-MS *computed for $\mathcal{M}_0(t)$. Suppose that for every $M \in \mathcal{M}(t) \setminus \mathcal{M}_0(t)$,*

$$\psi(t) \sum_{l \in M} w_l(t) q_l(t) \bar{\phi}_l(t) \le \max_{M' \in \mathcal{M}_0(t)} \sum_{l \in M'} q_l(t) w_l(t) \bar{\phi}_l(t) + (1-\psi(t)) \mathbb{E}^{\text{R,F}}\left[\mathcal{G}_{ALG}(t) | \mathcal{S}(t)\right], \tag{19}$$

for some $\psi(t) \in (0,1]$, where $\phi_l(t)$ was defined in (11) and $\bar{\phi}_l(t) = 1 - \phi_l(t)$. Then

$$\gamma^\star_{\text{ALG}} \ge \min_{t \ge 0} \Big\{\psi(t) \frac{|\mathcal{M}_0(t)|}{2|\mathcal{M}_0(t)|-1}\Big\}. \tag{20}$$

Note that we are interested in satisfying Condition (19) with the highest value of the parameter $\psi(t)$ to obtain a better efficiency ratio.

Next, by focusing on a special case in which Condition (19) holds, we can design provably efficient policies by considering a small set of schedules such that any other maximal schedule can be covered by them.

Lemma 4. *Suppose every $M \in \mathcal{M}(t) \setminus \mathcal{M}_0(t)$ is covered by at most ζ maximal schedules from $\mathcal{M}_0(t)$, i.e.,*

$$M \subseteq \cup_{M' \in S_M} M', \text{ for some } S_M \subseteq \mathcal{M}_0(t): \ |S_M| \leq \zeta. \quad (21)$$

Then Condition (19) holds with fixed $\psi(t) = \frac{1}{\zeta}$.

Proof. Using the covering definition (21), we have

$$\sum_{l \in M} q_l(t) w_l(t) \bar{\phi}_l(t) \leq \sum_{M' \in S_M} \sum_{l \in M'} q_l(t) w_l(t) \bar{\phi}_l(t)$$
$$\leq \zeta \max_{M' \in \mathcal{M}_0(t)} \sum_{l \in M'} q_l(t) w_l(t) \bar{\phi}_l(t), \quad (22)$$

and hence condition (19) trivially holds for $\psi = \frac{1}{\zeta}$. □

In general, $\mathcal{M}_0(t)$ can be adaptive and constructed based on the link deficits and fading probabilities. Here, we apply Proposition 1 and Lemma 4 for a constant $\psi(t) = \psi$ and a family $\mathcal{M}_0(t)$ induced by fixed subset of the independent sets $\mathcal{I}_0 \subseteq \mathcal{I}$, i.e., $\mathcal{M}_0(t) = \{D \cap \mathcal{B}(t), D \in \mathcal{I}_0\}$. With minor abuse of notation, we refer to such algorithms as FAMIX-MS$|_{\mathcal{I}_0}$. Below, we present a covering-based algorithm.

Corollary 4.1 (Coloring-based Randomization)**.** *Consider a coloring of graph G_I with χ colors, which partitions the vertices of G_I into χ independent sets $\{D'_1, ..., D'_\chi\}$. Extend these independent sets arbitrarily so they are maximal $\{D_1, ..., D_\chi\} := \mathcal{I}_0$. Then $\gamma^\star_{\text{FAMIX-MS}|_{\mathcal{I}_0}} \geq \frac{1}{2\chi - 1}$.*

Proof. Any maximal schedule in $\mathcal{M} \setminus \mathcal{M}_0$ can be covered by at most χ maximal schedules in $\mathcal{M}_0$, thus $\psi = \frac{1}{\chi}$ by Lemma 4. Further, $\frac{|\mathcal{M}_0(t)|}{2|\mathcal{M}_0(t)|-1} \geq \frac{\chi}{2\chi-1}$. Thus applying Proposition 1, we get the stated efficiency ratio. □

Remark 5. In general, finding an efficient coloring might be computationally demanding, but it needs to be done only once for a given G_I. There are many interesting families of graphs for which coloring can be solved efficiently. For example, for a (not necessarily complete) bipartite graph (e.g. a tree) where $\chi = 2$, we obtain $\gamma^\star \geq \frac{1}{3}$. This performs much better than LDF whose efficiency ratio in a tree with maximum degree β is $\gamma^\star \geq \frac{1}{2\sqrt{\beta - 1} + 1}$ in the case without fading (which is a special case of our setting). Another family that admits an efficient coloring are planar graphs where, by the four-color theorem, always have a 4-coloring which can be found in polynomial time [25]. Further, we remark that the independent sets D'_i could be extended adaptively at every time t, e.g., using GMS.

2) Myopic Distributed Randomized Algorithm: We present a simple distributed algorithm that has constant complexity.

Assume each slot is divided in two parts, a control part of duration T_C, and a packet transmission part with duration normalized to 1. At the beginning of the control phase, every non-empty link $l \in \mathcal{B}(t)$ starts a timer $T_l \sim Exp(\nu_l)$, where $Exp(\nu)$ denotes an exponential distribution with rate ν. Once the timer of a link l runs down to zero, it broadcasts an announcement informing its neighbors that it will participate in data transmission, *unless* it has heard an earlier announcement from its neighboring links, or the control phase ends.

Given any $\delta \in (0, 1)$, let $T_C = \max_{\nu_l}\{\frac{-\log(\delta)}{\nu_l}\}$. The next corollary states the efficiency ratio for the uniform timer rates.

Corollary 4.2. *Consider the myopic randomized algorithm where every link l has the same timer rate $\nu_l = \nu$. If the maximum degree of G_I is Δ, then*

$$\gamma^\star_{\mathit{MYOPIC}} \geq \frac{1}{\frac{\Delta - \delta}{1 - \delta} + 1} \quad (23)$$

Note that theoretically we can scale up the timer rate ν, so that the control phase T_C becomes very small.

Remark 6. We note that direct application of Theorem 1 yields an efficiency ratio $\approx \frac{1}{\Delta+2}$ as the myopic algorithm obtains $\approx \frac{1}{\Delta+1}$ approximation of the MWS (to see this note that every link has probability $\approx 1/(\Delta + 1)$ of getting service. Thus all the links of the MWS are included with this probability). Therefore Corollary 4.2 also serves as an example in which a careful analysis can improve the bound of direct application of Theorem 1.

IV. Analysis Techniques and Proofs

We provide an overview of the techniques in our proofs.

Frame Construction. A key step in the analysis of our scheduling algorithms is a frame construction similar to the one in [19], but based on the joint traffic-fading process. The definition of frame is as follows

Definition 1 (Frames and Cycles)**.** *Starting from an initial traffic and fading state tuple $(\boldsymbol{\tau}(0), \mathbf{q}(0)) = \mathbf{z} \in \mathcal{Z}$, let t_i denote the i-th return time of traffic-fading Markov chain $\mathbf{z}(t)$ to $\mathbf{z}$, $i = 1, \cdots$. 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 $\mathcal{C}_{(i-1)k+1}, \cdots, \mathcal{C}_{ik}$, 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 $(\boldsymbol{\tau}(t), \mathbf{q}(t))$ patterns during a frame $\mathcal{F}^{(k)}$. Note that these patterns start after $\mathbf{z}$ and end with $\mathbf{z}$.*

By the strong Markov property and the positive recurrence of traffic-fading Markov chain $\mathbf{z}(t)$, 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 $\mathcal{Z}$ 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-fading 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, $\overline{a}_l\mathbb{E}[F]=\overline{a}_l k\mathbb{E}[C]$, as $k\to\infty$. Define the rate region

$$\Lambda_{\rm NC}(\mathcal{F})=\bigcup_{\mu\in\mathcal{P}_{NC}(\mathcal{F})}\Lambda_\mu. \tag{24}$$

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

$$\lim_{t\to\infty}\frac{\sum_{s=1}^{t}I_l(s)}{t}=\frac{\mathbb{E}\left[\sum_{t\in\mathcal{F}}I_l(t)\right]}{\mathbb{E}[F]}=\bar{I}_l. \tag{25}$$

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

$$\frac{\mathbb{E}[\sum_{t\in\mathcal{F}}\widetilde{a}_l(t)]}{\mathbb{E}[F]}=\lambda_l,\ l\in\mathcal{K}. \tag{26}$$

In Definition 1, each frame consists of k cycles. Using similar arguments as in [16], [19], it is easy to see that

$$\liminf_{k\to\infty}\Lambda_{\rm NC}(\mathcal{F}^{(k)})\supseteq{\rm int}(\Lambda),$$

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

$$\rho\,{\rm int}(\Lambda_{\rm NC}(\mathcal{F}^{(k)}))\subseteq\Lambda_{\rm ALG}, \tag{27}$$

then it follows that $\Lambda_{\rm ALG}\supseteq\rho\,{\rm int}(\Lambda)$. For our algorithms, we find a ρ such that (27) holds for *any traffic-fading process* under our model. Then it follows that $\gamma^\star_{\rm ALG}\geq\rho$.

Lyapunov Argument. To prove (27), we rely on comparing the expected gain of ALG with that of the non-causal policy that maximizes the expected gain over the frame (max-gain policy). The following proposition, which is similar to that in [19], will be used to prove the main results. We omit its proof, as it is similar to the proof in [19] with minor modifications to account for channel uncertainty.

Proposition 2. *Consider a frame $\mathcal{F}\equiv\mathcal{F}^{(k)}$, for a fixed k based on the returns of the traffic-fading process $\mathbf{z}(t)$ to a state $\mathbf{z}$. Define the norm of initial deficits at the beginning of a frame $\|\mathbf{w}(t_0)\|=\sum_{l\in\mathcal{K}}w_l(t_0)$. Suppose for a causal policy* ALG*, given any $\epsilon>0$, there is a W' such that when $\|\mathbf{w}(t_0)\|>W'$,*

$$\frac{\mathbb{E}\left[\sum_{t\in\mathcal{F}}\mathcal{G}_{ALG}(t)|\mathcal{S}(t_0)\right]}{\mathbb{E}\left[\sum_{t\in\mathcal{F}}\mathcal{G}_{\mu^\star}(t)|\mathcal{S}(t_0)\right]}\geq\rho-\epsilon, \tag{28}$$

where $\mathcal{S}(t_0)=(\mathbf{q}(t_0),\boldsymbol{\tau}(t_0),\Psi(t_0),\mathbf{w}(t_0))$, and $\mu^\star$ is the non-causal policy that maximizes the gain over the frame. Then for any $\lambda\in\rho\,{\rm int}(\Lambda_{NC}(\mathcal{F}))$, the deficit queues are bounded in the sense of (6).

Amortized Gain Analysis. To use Proposition 2, we need to analyze the achievable gain of ALG and the non-causal policy $\mu^\star$ over a frame. Since comparing the gains of the two policies directly is difficult, we adapt an amortized analysis technique from [19], initially extended from [27]–[30]. The general idea is as follows. Let $(\mathbf{q}(t),\boldsymbol{\tau}(t),\Psi(t),\mathbf{w}(t))$ be the state under our algorithm at time $t\in\mathcal{F}$, and $(\Psi^{\mu^\star}(t),\mathbf{q}(t),\boldsymbol{\tau}(t),\mathbf{w}^{\mu^\star}(t))$ be the state under the optimal policy $\mu^\star$. The traffic-fading process $\mathbf{z}(t)=(\mathbf{q}(t),\boldsymbol{\tau}(t))$ is identical for both algorithms as it is independent of the actions of the scheduling policy. We change the state of $\mu^\star$ (by modifying its buffers and deficits) to make it identical to $(\mathbf{q}(t),\boldsymbol{\tau}(t),\Psi(t),\mathbf{w}(t))$, but also give $\mu^\star$ an additional gain that ensures the change is advantageous for $\mu^\star$ considering the rest of the frame. Let $\hat{\mathcal{G}}_{\mu^\star}(t)$ denote the amortized gain of $\mu^\star$ at time t with any compensated gain, which has the property that

$$\mathbb{E}\Big[\sum_{t\in\mathcal{F}}\hat{\mathcal{G}}_{\mu^\star}(t)|J,\mathcal{S}(t_0)\Big]\geq\mathbb{E}\Big[\sum_{t\in\mathcal{F}}\mathcal{G}_{\mu^\star}(t)|J,\mathcal{S}(t_0)\Big], \tag{29}$$

given any traffic-fading pattern $J\in\mathcal{J}(\mathcal{F})$ and initial frame state $\mathcal{S}(t_0)$. Then, the following proposition will be useful in bounding the gain and thus the efficiency ratio of our policies.

Proposition 3. *Consider a Markov policy* ALG *that for any traffic-fading pattern $J\in\mathcal{J}(\mathcal{F})$, at any time $t\in\mathcal{F}$, satisfies*

$$\rho\mathbb{E}\left[\hat{\mathcal{G}}_{\mu^\star}(t)|J,\mathcal{S}(t)\right]\leq\mathbb{E}\left[\mathcal{G}_{ALG}(t)|J,\mathcal{S}(t)\right]+\mathcal{E}_F \tag{30}$$

for some $\mathcal{E}_F$ which is a measurable function of the frame length F, with $\mathbb{E}\left[F\mathcal{E}_F\right]<\infty$. Then $\gamma^\star_{\rm ALG}\geq\rho$.

Proof. We mention a sketch of the proof but more detailed analysis can be found in the technical report [23]. First, using the Markov property of ALG, the fact that the amortized gain of $\mu^\star$ does not depend on the past state given the current state and future traffic-fading pattern, and summing over the frame, we can show

$$\rho\mathbb{E}\Big[\sum_{t\in\mathcal{F}}\hat{\mathcal{G}}_{\mu^\star}(t)|J,\mathcal{S}(t_0)\Big]\leq\mathbb{E}\Big[\sum_{t\in\mathcal{F}}\mathcal{G}_{ALG}(t)|J,\mathcal{S}(t_0)\Big]+F\mathcal{E}_F.$$

Second, using definition (29), and taking expectation w.r.t. traffic-fading pattern, we have

$$\rho\mathbb{E}\Big[\sum_{t\in\mathcal{F}}\mathcal{G}_{\mu^\star}(t)|\mathcal{S}(t_0)\Big]\leq\mathbb{E}\Big[\sum_{t\in\mathcal{F}}\mathcal{G}_{ALG}(t)|\mathcal{S}(t_0)\Big]+\mathbb{E}\left[F\mathcal{E}_F\right].$$

Third, $\mathbb{E}\left[F\mathcal{E}_F\right]<\infty$, and by the the non-trivial traffic-fading Markov chain assumption (Section II), we can show $\lim_{\|\mathbf{w}(t_0)\|\to\infty}\mathbb{E}\left[\sum_{t\in\mathcal{F}}\mathcal{G}_{\mu^\star}(t)|\mathcal{S}(t_0)\right]=\infty$. Then, we can apply Proposition 2 to get $\gamma^\star_{\rm ALG}\geq\rho$. □

The following lemma describes a generic amortized gain computation for general networks that allows us to modify the state of the max-gain policy during a frame to match the state of the considered policy ALG. Recall from (11) that $\phi_l(t)$ is the probability that link l is scheduled under ALG, and $\pi_M(t)$ is the probability that schedule M is selected. Below we drop their dependence on time to simplify the notation.

Lemma 5. *For any pattern $J\in\mathcal{J}(\mathcal{F})$ in a frame $\mathcal{F}$, given a Markov policy* ALG*, the amortized gain of the max-gain policy $\mu^\star$ (w.r.t.* ALG*) if it selects maximal schedule $M\in\mathcal{M}$ at time t in the frame, is given by*

$$\begin{aligned}\mathbb{E}^{\rm R,F}\left[\hat{\mathcal{G}}^{(M)}_{\mu^\star}(t)|J,\mathcal{S}(t)\right]&=W_M(t)+\sum_{l\notin M}w_l(t)\phi_l q_l(t)+\mathcal{E}_m\\&\leq W_M(t)(1-\pi_M)+\mathbb{E}^{\rm R,F}\left[\mathcal{G}_{ALG}(t)|\mathcal{S}(t)\right]+\mathcal{E}_m,\end{aligned}$$

where $\mathcal{E}_m:=K(F+a_{max}d_{\max})$.

Proof. The main idea is similar to the one in [19], but we have to account for fading. Suppose $\mu^\star$ attempts to transmit the earliest-deadline packets of links from schedule M whereas ALG attempts the earliest-deadline packets from schedule M'. We need to modify the state of $\mu^\star$, i.e., the buffers and deficits, so it is identical with the state of ALG. To achieve this, we allow $\mu^\star$ to additionally transmit the packets of links successfully transmitted by ALG but not $\mu^\star$, i.e, $S_1 = (M'(t) \setminus M(t)) \cap \{l \in \mathcal{K} : C_l(t) = 1\}$. Transmitting such packets for $\mu^\star$ might not be advantageous for its total gain as the weight of these packets can increase by $a_{max} d_{\max}$ before they could be transmitted. Thus giving an additional reward $K a_{max} d_{\max}$ to $\mu^\star$ guarantees that the modification is advantageous. Further, we insert the packets transmitted successfully by $\mu^\star$ but not ALG, i.e., $S_2 = (M(t) \setminus M'(t)) \cap \{l \in \mathcal{K} : C_l(t) = 1\}$, back to its buffers (which is advantageous for $\mu^\star$). Further to make the deficits identical, we increase the deficit counters of $\mu^\star$ for links in S_2, which is advantageous for the total-gain within the frame. Additionally, we decrease the deficit counters of $\mu^\star$ for the links in S_1. This change might not be advantageous for the total gain of $\mu^\star$, thus we give it extra reward for every possible subsequent transmission over links, which is at most KF. Hence the total additional compensation is $\mathcal{E}_m = K(F + a_{max} d_{\max})$

Following the above argument, the expected amortized gain of $\mu^\star$, when it selects M, is bounded as

$$\mathbb{E}^{\mathrm{R,F}}\left[\hat{\mathcal{G}}^{(M)}_{\mu^\star}(t)|J, \mathcal{S}(t)\right] = \sum_{l\in M} q_l(t) w_l(t) + \sum_{M'\in\mathcal{M}\setminus\{M\}} \pi_{M'}(t) \sum_{l\in M'\setminus M} q_l(t) w_l(t) + \mathcal{E}_m$$
$$\stackrel{(a)}{=} W_M(t) + \sum_{l\notin M} w_l(t)\phi_l q_l(t) + \mathcal{E}_m \quad (31)$$

where (a) follows from definitions of ϕ_l and $W_M(t)$. The inequality in the lemma's statement follows by noting that

$$(31) \leq W_M(t) + \sum_{M'\in\mathcal{M}\setminus\{M\}} W_{M'}(t)\pi_{M'} + \mathcal{E}_m,$$

since $\sum_{l\in M'\setminus M} q_l(t) w_l(t) \leq W_{M'}(t)$, and by using (12). □

A. Proof of Theorem 1: Conversion Result

In the rest of proofs, for notional compactness, we define $\mathbb{E}_{t,J}[\cdot] := \mathbb{E}[\cdot|\mathcal{S}(t), J]$ and $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot|\mathcal{S}(t)]$. Now consider a pattern $J \in \mathcal{J}(\mathcal{F})$. When $\|\mathbf{w}(t)\| \geq W'$, the amortized gain of $\mu^\star$, if it selects schedule M, is bounded as

$$\begin{aligned}\mathbb{E}^{\mathrm{R,F}}_{t,J}\left[\hat{\mathcal{G}}^{(M)}_{\mu^\star}(t)\right] &\stackrel{(a)}{\leq} W_M(t)(1-\pi_M) + \mathbb{E}^{\mathrm{R,F}}_t\left[\mathcal{G}_{ALG}(t)\right] + \mathcal{E}_m \\ &\leq W_{M^\star}(t) + \mathbb{E}^{\mathrm{R,F}}_t\left[\mathcal{G}_{ALG}(t)\right] + \mathcal{E}_m \\ &= \mathbb{E}^{\mathrm{R,F}}_t\left[\mathcal{G}_{ALG}(t)\right]\left(\frac{W_{M^*}(t)}{\mathbb{E}^{\mathrm{R,F}}_t\left[\mathcal{G}_{ALG}(t)\right]} + 1\right) + \mathcal{E}_m \\ &\stackrel{(b)}{\leq} \mathbb{E}^{\mathrm{R,F}}_t\left[\mathcal{G}_{ALG}\right](1/\psi + 1) + \mathcal{E}_m,\end{aligned}$$

where in (a) we used Lemma 5, and in (b) we used the main assumption that ALG obtains ψ fraction of the maximum weight schedule. This inequality does not depend on the particular choice M. Similarly in the case that $\|\mathbf{w}(t)\| \leq W'$, it can be seen from (13) and Lemma 5 that $\mathbb{E}^{\mathrm{R,F}}_{t,J}\left[\hat{\mathcal{G}}^{(M)}_{\mu^\star}(t)\right] \leq 2W' + \mathcal{E}_m$. Consequently in either case, $\mathbb{E}^{\mathrm{R,F}}_{t,J}\left[\hat{\mathcal{G}}^{(M)}_{\mu^\star}(t)\right] \leq 2W' + \mathcal{E}_m + \mathbb{E}^{\mathrm{R,F}}_t\left[\mathcal{G}_{ALG}(t)\right](1/\psi + 1)$. Applying Proposition 3 with $\rho = (1/\psi + 1)^{-1}$, we obtain the result.

B. Analysis of FAMIX-MS*: Proof of Theorem 2.*

We use Lemma 5 and using probabilities of FAMIX-MS (16) we can show (30) and apply Proposition 3. Detailed derivation can be found in the technical report [23].

C. Analysis of FAMIX-ND*: Proof of Theorem 3*

We first state the following Lemma that allows us to focus on policies that transmit from non-dominated links.

Lemma 6. *Given $J \in \mathcal{J}(\mathcal{F})$, let $\tilde{\mu}^\star$ be the maximum-gain policy that transmits only from non-dominated links at any time (we refer to $\tilde{\mu}^\star$ as max-gain ND-policy) and the maximum-gain policy $\mu^\star$ that can transmit any packet, then*

$$\mathbb{E}\left[\textstyle\sum_{t\in\mathcal{F}} \mathcal{G}_{\tilde{\mu}^\star}(t)|\mathcal{S}(t_0), J\right] \geq \mathbb{E}\left[\textstyle\sum_{t\in\mathcal{F}} \mathcal{G}_{\mu^\star}(t)|\mathcal{S}(t_0), J\right] - (a_{max}+1)F^2.$$

Proof. The proof has two parts: (i) Consider the earliest-deadline packets $\mathtt{P}_u, \mathtt{P}_v$ of links u, v, with u dominating v. Assuming deficits are fixed, we can show through a coupling [31] that transmitting $\mathtt{P}_u$ yields a higher expected gain than that of $\mathtt{P}_v$ over the rest of the frame. (ii) The change of deficits over the frame is bounded by a multiple of the frame length F, which adds a correction term $(a_{max}+1)F^2$ to the expected gain. More details can be found in the technical report [23]. □

Proof of Theorem 3. In view of Lemma 6, we perform amortized analysis for FAMIX-ND in comparison with ND-policies. Suppose FAMIX-ND decides to schedule earliest-deadline packet $\mathtt{P}_f = (w_f, e_f)$ of link f and ND-policy $\tilde{\mu}^\star$ transmits packet $\mathtt{P}_z = (w_z, e_z)$ from link z. Depending on the outcome of the channels and the selected z, f, similarly to the arguments in the proof of Lemma 5, we modify the state of $\tilde{\mu}^\star$ and give appropriate compensation. We need to analyze a total of 4 cases. The difference from Lemma 5 lies in the case that $e_f \leq e_z, w_f \leq w_z$, and $C_f(t) = 1, C_z(t) = 1$. In this case, replacing $\mathtt{P}_f$ in link f with $\mathtt{P}_z$ in link z suffices to make the buffers of $\tilde{\mu}^\star$ identical to FAMIX-ND. The details of all the cases can be found in the technical report [23]. As a result of this analysis, we can show

$$\begin{aligned}\mathbb{E}_{t,J}\left[\hat{\mathcal{G}}^{(i)}_{\tilde{\mu}^\star}(t)\right] &\leq q w_i + \sum_j \pi_j (q - q_{ij}) w_j + \sum_{j<i} w_j \pi_j q_{ij} + \mathcal{E}_0 \\ &\leq q[w_i + (1-q)\sum_j \pi_j w_j + q\sum_{j<i} w_j \pi_j] + \mathcal{E}_0\end{aligned}$$

where $q_{ij} = \Pr(C_i(t) = 1, C_j(t) = 1)$, and $\mathcal{E}_0 = F + (a_{max} + 1)d_{\max}$. Note that since $\pi_i \leq \frac{1}{q}(1 - \frac{w_{i+1}}{w_i})$, the right-hand-side of the above inequality is maximized for $i = 1$. Hence,

$$\mathbb{E}_{t,J}\left[\hat{\mathcal{G}}_{\tilde{\mu}^\star}(t)\right] \leq q w_1 + (1-q)\mathbb{E}_t\left[\mathcal{G}_{ALG}(t)\right] + \mathcal{E}_0. \quad (32)$$

Let $\bar{n}$ be the number of links with positive probability in (17). For the gain of FAMIX-ND, we can obtain

$$\mathbb{E}_t\left[\mathcal{G}_{ALG}(t)\right] = \sum_{i=1}^{\bar{n}} q\pi_i w_i$$
$$\stackrel{(a)}{=} w_1\left(1-(1-q+q\sum_{i=1}^{\bar{n}-1}\pi_i)\prod_{i=1}^{\bar{n}-1}(1-\pi_i q)\right)$$
$$\stackrel{(b)}{\geq} w_1(1-\left(\frac{\bar{n}-q}{\bar{n}}\right)^{\bar{n}}) \geq w_1(1-e^{-q}),$$

where in (a), by Definition of π_i in (17) the sum telescopes and we have: $w_{\bar{n}} = w_1\prod_{i=1}^{\bar{n}-1}(1-\pi_i q)$, and (b) follows from the geometric-arithmetic inequality. Using the above relation and (32) we get

$$\mathbb{E}_{t,J}\left[\hat{\mathcal{G}}_{\tilde{\mu}^\star}(t)\right] \leq \mathbb{E}_t\left[\mathcal{G}_{ALG}(t)\right](\frac{q}{1-e^{-q}}+(1-q))+\mathcal{E}_0.$$

Summing over the entire frame using similar arguments as in proof of Proposition 3, as well as using Lemma 6, it follows:

$$\mathbb{E}\left[\sum_{t\in\mathcal{F}}\mathcal{G}_{\mu^\star}(t)|J,\mathcal{S}(t_0)\right]-F^2(a_{max}+1)\leq$$
$$\mathbb{E}\left[\sum_{t\in\mathcal{F}}\mathcal{G}_{ALG}(t)|J,\mathcal{S}(t_0)\right](\tfrac{q}{1-e^{-q}}+1-q)+F\mathcal{E}_0.$$

The proof is concluded similarly as in the proof of Proposition 3, applying Proposition 2 with $\rho = (\frac{q}{1-e^{-q}}+1-q)^{-1}$.

D. Proof of Proposition 1: Low-Complexity Variants

Suppose that under ALG we have

$$\psi(t)\max_{M\in\mathcal{M}}\mathbb{E}_{t,J}\left[\hat{\mathcal{G}}^{(M)}_{\mu^\star}(t)\right] \leq \max_{M\in\mathcal{M}_0(t)}\mathbb{E}_{t,J}\left[\hat{\mathcal{G}}^{(M)}_{\mu^\star}(t)\right]. \quad (33)$$

Since ALG randomizes according to the probabilities of FAMIX-MS over the maximal schedules in $\mathcal{M}_0(t)$, we have

$$\xi(t)\max_{M\in\mathcal{M}_0(t)}\{\mathbb{E}_{t,J}\left[\hat{\mathcal{G}}^{(M)}_{\mu^\star}(t)\right]\} \leq \mathbb{E}_t\left[\mathcal{G}_{ALG}(t)\right]+\mathcal{E}, \quad (34)$$

where $\xi(t) := \frac{|\mathcal{M}_0(t)|}{2|\mathcal{M}_0(t)|-1}$, by using the same arguments as in the proof of Theorem 2 applied to $\mathcal{M}_0(t)$. By (33) and (34),

$$\min_{t\geq 0}\{\xi(t)\psi(t)\}\,\mathbb{E}_{t,J}\left[\hat{\mathcal{G}}_{\mu^\star}(t)\right] \leq \mathbb{E}_t\left[\mathcal{G}_{ALG}(t)\right]+\mathcal{E}.$$

Then by Proposition 3, it follows that $\gamma \geq \min_{t\geq 0}\{\psi(t)\xi(t)\}$. Hence to conclude the proof, we simply need to argue that (19) implies (33) Using Lemma 5, and (13), it can be shown that

$$\mathbb{E}^{\mathrm{R,F}}_{t,J}\left[\hat{\mathcal{G}}^{(M)}_{\mu^\star}(t)\right] = \sum_{l\in M} w_l(t)q_l(t)\bar{\phi}_l + \mathbb{E}^{\mathrm{R,F}}_t\left[\mathcal{G}_{ALG}(t)\right]+\mathcal{E}_m$$

Plugging the above expression in (33), it can be seen that (19) is a sufficient condition for (33) to hold.

E. Proof of Corollary 4.2: Myopic Distributed Algorithm

First, we can show that under the myopic randomized algorithm, $\phi_l \geq \frac{1-\delta}{1+\Delta}$ for every nonempty link. Using this result, we can do the gain analysis based on Lemma 5, and (13), and show that Proposition 3 holds for $\rho = (\frac{\Delta-\delta}{1-\delta}+1)^{-1}$. For more details refer to the technical report [23].

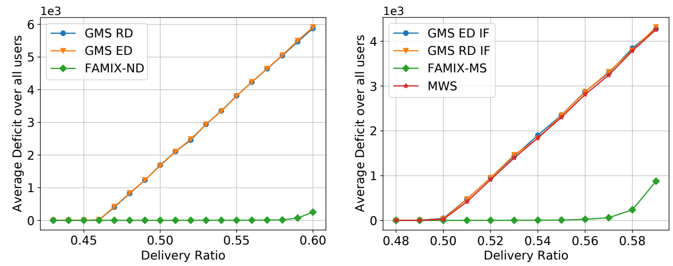

(a) 2-link collocated network.   (b) 5-link sparse network

Fig. 2: Performance comparison of various algorithms.

V. Simulation Results

We performed simulations under several networks and traffic-fading scenarios. Our algorithms FAMIX-ND and FAMIX-MS can considerably outperform GMS and MWS. We present two of the simulations here, but more simulations can be found in the technical report [23].

First, we consider the traffic of Pattern B of Figure 1, but with i.i.d. channel success probability q. The results are shown in Figure 2a for $q = 0.6$ and equal target delivery ratio p for links. Note that the optimal cannot achieve $p > 0.6$, hence FAMIX-ND is near optimal. We observed a similar behavior for other values of q and different number of links.

Next, consider a network with 5 links and G_I with edges $\{(l_1,l_2),(l_2,l_3),(l_2,l_4),(l_4,l_5)\}$. The traffic-fading process for $\{l_1,l_3,l_4\}$ is as in link 2 of Pattern B in Figure 1, and for links $\{l_2,l_5\}$ is as in link 1. We set an equal target delivery ratio p for all the links. Figure 2b shows the results. We see FAMIX-MS can support a significantly higher delivery ratio than other algorithms. In this case, optimal cannot achieve $p > 0.75$.

Next, consider a simple network with links $\mathcal{K} = \{l_1,l_2,l_3\}$ and edges $\{(l_1,l_2),(l_1,l_3)\}$ in the interference graph G_I with Pattern A from Figure 1. Then GMS becomes unstable for $\mathbf{p} = (\frac{1}{3}+\epsilon,\frac{1}{2}+\epsilon,\frac{1}{2}+\epsilon)$, whereas the optimal can satisfy $\mathbf{p} = (\frac{11}{12},\frac{5}{6},\frac{5}{6})$. This example illustrates the existence of simple cases with non-equal delivery ratio requirements between users where GMS performs poorly. More details can be found in the technical report [23]. Similar behavior was observed for examples with more than three links.

VI. Conclusion

We considered scheduling of real-time traffic over fading channels, where traffic (arrival and deadline) and links' reliability evolves as an *unknown* finite-state Markov chain. We provided a conversion result that shows classical non-real-time scheduling algorithms like MWS and GMS can be ported to this setting and characterized their efficiency ratio. We then extended the randomized algorithms from [19] to fading channels. We further proposed low-complexity and myopic distributed randomized algorithms and characterized their efficiency ratio. Investigating more efficient low-complexity and distributed randomized algorithms could be an interesting future research.

References

[1] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," in *29th IEEE Conference on Decision and Control*, 1990, pp. 2130–2132.

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] J. Song, S. Han, A. Mok, D. Chen, M. Lucas, M. Nixon, and W. Pratt, "Wirelesshart: 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.

[9] 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.

[10] 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.

[11] 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.

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

[13] 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.

[14] B. Li and A. Eryilmaz, "Optimal distributed scheduling under time-varying conditions: A fast-csma algorithm with applications," *IEEE Transactions on Wireless Communications*, vol. 12, no. 7, pp. 3278–3288, 2013.

[15] R. Singh and P. Kumar, "Throughput optimal decentralized scheduling of multihop networks with end-to-end deadline constraints: Unreliable links," *IEEE Transactions on Automatic Control*, vol. 64, no. 1, pp. 127–142, 2018.

[16] 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.

[17] 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.

[18] L. Deng, C.-C. Wang, M. Chen, and S. Zhao, "Timely wireless flows with general traffic patterns: Capacity region and scheduling algorithms," *IEEE/ACM Transactions on Networking*, vol. 25, no. 6, pp. 3473–3486, 2017.

[19] C. Tsanikidis and J. Ghaderi, "On the power of randomization for scheduling real-time traffic in wireless networks," in *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, 2020, pp. 59–68.

[20] I.-H. Hou, "Scheduling heterogeneous real-time traffic over fading wireless channels," *IEEE/ACM Transactions on Networking*, vol. 22, no. 5, pp. 1631–1644, 2013.

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

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

[23] C. Tsanikidis and J. Ghaderi, "Randomized scheduling of real-time traffic in wireless networks over fading channels," *arXiv preprint arXiv:2101.04815*, 2021.

[24] C. Joo, X. Lin, J. Ryu, and N. B. Shroff, "Distributed greedy approximation to maximum weighted independent set for scheduling with fading channels," *IEEE/ACM Transactions on Networking*, vol. 24, no. 3, pp. 1476–1488, 2015.

[25] N. Robertson, D. Sanders, P. Seymour, and R. Thomas, "A new proof of the four-colour theorem," *Electronic Research Announcements of the American Mathematical Society*, vol. 2, no. 1, pp. 17–25, 1996.

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

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

[28] 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.

[29] 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.

[30] 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.

[31] T. Lindvall, *Lectures on the coupling method*. Courier Corporation, 2002.