

On the Performance of Largest-Deficit-First for Scheduling Real-Time Traffic in Wireless Networks

Xiaohan Kang, Weina Wang, *Student Member, IEEE*, Juan José Jaramillo, *Member, IEEE*,
and Lei Ying, *Member, IEEE*

Abstract—This paper considers the problem of scheduling real-time traffic in wireless networks. We consider ad hoc wireless networks with general conflict graph-based interference model and single-hop traffic. Each packet is associated with a deadline and will be dropped if it is not transmitted before the deadline. The number of packet arrivals in each time slot and the maximum delay before the deadline are independent and identically distributed across time. We require a minimum fraction of packets to be delivered. At each link, we assume the link keeps track of the difference between the minimum number of packets that need to be delivered so far and the number of packets that are actually delivered, which we call the deficit. The largest-deficit-first (LDF) policy schedules links in descending order according to their deficit values, which is a variation of the longest-queue-first (LQF) policy for non-real-time traffic. We prove that the efficiency ratio of LDF, which is the fraction of the throughput region that LDF can achieve for given traffic distributions, can be lower bounded by a quantity that we call the real-time local-pooling factor (R-LPF). We further prove that a lower bound on the R-LPF can be related to the weighted sum of the service rates, with a special case of $1/(\beta + 1)$ by considering the uniform weight, where β is the interference degree of the conflict graph. We also propose a heuristic consensus algorithm that can be used to obtain a good weight vector for such lower bounds for given network topology.

Index Terms—Stability, real-time scheduling, largest-deficit-first, local-pooling factor, fluid limit.

I. INTRODUCTION

With the increasing number of real-time applications in wireless networks, scheduling traffic of packets with hard deadlines has become a very important problem. However, the problem is very challenging due to the stochastic nature of the traffic arrivals and deadlines. Hou et al. first proposed a frame-based analytical framework for studying scheduling real-time traffic in wireless networks [1]. In the frame-based framework it is assumed that each frame is a number of consecutive time slots, and all packets arrive at the beginning of a frame and have to be scheduled before the end of the frame. They also characterized the real-time capacity region and developed the optimal scheduling algorithm for collocated networks. Later, the frame-based framework has been generalized to networks with heterogeneous delays, fading, congestion control, etc. [2]–[6] In particular, Jaramillo et al. extended the idea to general arrival/deadline patterns within a

frame and general non-collocated network topology, and found the optimal scheduling policy [6], where they assumed that packets can arrive at any time slot during a frame, and the deadline of a packet can be any time after its arrival and before the end of the frame. Their paper assumes that the arrival and deadline information is available at the beginning of the frame, so future knowledge is assumed. Furthermore, the computational complexity of the optimal algorithm is prohibitively high except for some special cases such as collocated networks.

In this paper, we consider the case of general real-time traffic patterns without the assumption of frames and with a general conflict graph-based interference model. Under these settings, the stability region is difficult to characterize, and the optimal policy is unknown. In this paper, we are interested in the performance of a low-complexity greedy policy called the largest-deficit-first (LDF) policy [1], which is the real-time variation of the longest-queue-first (LQF) policy that iteratively selects the link with the largest deficit that does not interfere with those links that are already selected. It has been shown that the largest-deficit-first policy is optimal for scheduling real-time traffic in collocated networks [1], [6] under the frame-based model. The performance of the LDF in general non-collocated networks has not been studied.

Since LDF can be directly applied to networks with non-frame-based real-time traffic, we are interested in characterizing the performance of LDF. We investigate the efficiency ratio of LDF, which is the fraction of the throughput region guaranteed by LDF for given traffic distributions. Although the capacity region and optimal scheduling algorithm for networks with non-frame-based real-time traffic remain unknown, we are able to establish the efficiency ratio of LDF by connecting it to the frame-based optimal scheduling algorithm, and obtain a lower bound on the efficiency ratio in terms of a new quantity, called the real-time local-pooling factor (R-LPF). The R-LPF extends the idea of the local-pooling factor for non-real-time traffic [7] and its extension for fading channels [8].

We show using the fluid limit technique [9] that this R-LPF can be successfully used to provide a minimum performance guarantee of LDF under real-time traffic. While the R-LPF depends on the traffic pattern, we lower-bound the R-LPF by purely topological quantities based on the network, which in particular connects the R-LPF with the interference degree of the conflict graph¹ [10]. Our contributions are therefore

X. Kang, W. Wang and L. Ying are with the School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ 85281 USA e-mail: {xiaohan.kang, weina.wang, lei.ying.2}@asu.edu.

J. J. Jaramillo is with the Department of Applied Math and Engineering, Universidad EAFIT, Medellín, Colombia e-mail: jjaram93@eafit.edu.co.

¹The interference degree of a network with a conflict graph is the maximum number of links that interfere with some single link and can be scheduled simultaneously.

fourfold:

- 1) We formulate the construction of the R-LPF and prove that it is a lower bound on the efficiency ratio of LDF in the presence of non-frame-based real-time traffic.
- 2) We show that by assigning a nonnegative weight to each link we can get a lower bound on the R-LPF regardless of the traffic pattern. This translates to an R-LPF at least $1/(\beta + 1)$ for a network with interference degree β , and in particular an R-LPF at least $1/3$ in a network with one-hop interference model.
- 3) We also propose a heuristic consensus algorithm that intelligently assigns the weights to compute a good lower bound based on the network topology.
- 4) We evaluate the performance of the LDF policy and the proposed consensus algorithm via simulations.

We would like to emphasize again that for general (non-frame-based) real-time traffic, to the best of our knowledge, there are no known theoretical results on any scheduling policy in ad hoc wireless networks. This makes the lower bounds obtained in this paper a novel contribution.

II. MODEL

In this paper, we consider a wireless network consisting of K links. The set of links is denoted by \mathcal{K} . Assume time is slotted, and at each time slot one packet can be successfully transmitted over a link if no interfering links are transmitting at the same time. We remark that the constant service rate assumption has been widely used in the literature, e.g., [11], [12]. We consider a general interference model. We call a set of links $\mathcal{Z} \subseteq \mathcal{K}$ a maximal link schedule if links in \mathcal{Z} can be scheduled at the same time without interfering with each other, but no other link can be further scheduled without interfering with links in \mathcal{Z} . We assume that there are R possible maximal link schedules and the set of maximal link schedules is represented by a maximal link schedule matrix M , which is a K -by- R matrix with binary entries such that each column represents a distinct maximal link schedule and the set of links that are included in this schedule have value 1 in that column. For example, let M_r be the r th column of matrix M , then the set of links $\{l \in \mathcal{K} : M_{l,r} = 1\}$ is a maximal link schedule, where $M_{l,r}$ is the (l, r) entry of the matrix. By abuse of notation we also let $M = \{M_1, M_2, \dots, M_R\}$. It is easy to see that any subset of a maximal link schedule is itself a feasible link schedule (i.e., all links in that set can be scheduled at the same time).

We consider single hop traffic with deadline constraints. Let $a_l(t)$ denote the number of packets that arrive at the beginning of time slot t at link l , where we assume that all packets have the same size and can be transmitted in a single time slot. We assume that $\{a_l(t), t \geq 1\}$ is a stochastic process that is temporally independent and identically distributed (i.i.d.) and independent across links, with probability mass function (p.m.f.) $f_l: \mathbb{N} \rightarrow \mathbb{R}$, where \mathbb{N} is the set of nonnegative integers and \mathbb{R} is the set of real numbers. We also assume that $f_l(i) = 0$ for $i > a_{\max}$; i.e., the number of packets arriving on a link at each time slot is at most a_{\max} . Denote by \bar{a}_l the arrival rate on link l ; i.e., $\bar{a}_l = \mathbb{E}[a_l(t)] = \sum_{i=1}^{a_{\max}} i f_l(i)$ for any t .

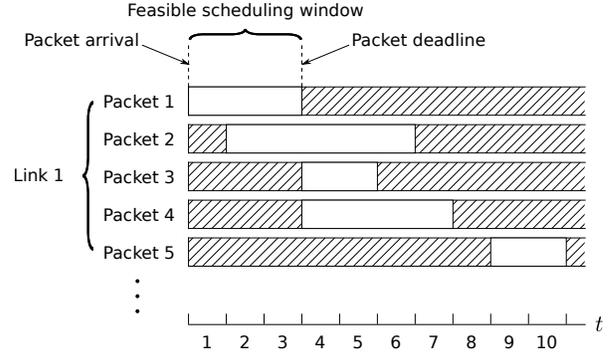


Fig. 1: An example of the arrival and maximum delay pattern of packets on a link. For each packet, the beginning of the blank bar is the time slot when that packet arrives, and the end of the blank bar is the deadline associated with that packet. So the feasible scheduling window denoted by the blank bar represents the time slots when the packet is available for transmission, while the shaded part indicates that the packet is not available, either because it has not arrived or because its deadline has passed. Note that here the cumulative numbers of packet arrivals to link 1 are $\{A_1(t), t \geq 0\} = (0, 1, 2, 2, 4, 4, 4, 4, 5, 5, \dots)$ and the maximum delays are $\{\tau_1(n), n \geq 1\} = (3, 5, 2, 4, 2, \dots)$.

Each packet is associated with a maximum delay τ , which is a random variable with integer value between τ_{\min} and τ_{\max} and follows a p.m.f. $g_l: \{\tau_{\min}, \tau_{\min} + 1, \dots, \tau_{\max}\} \rightarrow \mathbb{R}$. Furthermore, let $A_l(t)$ be the cumulative number of packet arrivals to link l up to time slot t for any $l \in \mathcal{K}$ and any nonnegative integer t ; i.e., $A_l(t) = \sum_{t'=1}^t a_l(t')$, and by convention $A_l(0) = 0$. We order the packets arriving on link l according to the arriving time with arbitrary tie-breakings. Then we let $b_l(n)$ be the time slot during which the n th packet arrives on link l ; i.e., $b_l(n) = \min\{t : A_l(t) \geq n\}$. We also let $e_l(n)$ be the deadline of the n th packet on link l . Note that $e_l(n) = b_l(n) + \tau_l(n) - 1$, where $\tau_l(n)$ is the maximum delay associated with the n th packet on link l . Then the n th arriving packet on link l will be immediately dropped if the deadline is missed. Note that $\{A_l(t), t \geq 0\}$, $\{\tau_l(n), n \geq 1\}$, $\{b_l(n), n \geq 1\}$ and $\{e_l(n), n \geq 1\}$ are all stochastic processes, and $\{A_l(t), t \geq 0\}$ and $\{\tau_l(n), n \geq 1\}$ determine $\{b_l(n), n \geq 1\}$ and $\{e_l(n), n \geq 1\}$. Denote the space of sample paths of the cumulative arrival process $\{A_l(t), t \geq 0\}$ and the maximum delay process $\{\tau_l(n), n \geq 1\}$ by \mathcal{A} . An example of a sample path of the arrival and maximum delay processes on a link during the first 10 time slots is shown in Fig. 1.

We assume each link l in \mathcal{K} is associated with a minimum delivery rate p_l (sometimes called the *quality of service* or *QoS*), which is the minimum fraction of packets that should be delivered on link l . The goal of a scheduling policy is to keep the long-term delivery rate on link l at least p_l .

Now consider a scheduling policy μ . Denote by $S^\mu(t)$ the cumulative service up to time t , in which $S_l^\mu(t)$ is the service link l received up to time slot t . For any scheduling policy, it is easy to see that the following three conditions hold:

- 1) (Initialization) $S_l^\mu(0) = 0$ for all $l \in \mathcal{K}$.

- 2) (Feasibility) The incremental service vector is a feasible schedule; i.e., $0 \preceq S^\mu(t) - S^\mu(t-1) \preceq M_r$ for some $M_r \in M$, for any positive integer t , where \preceq denotes entrywise less than or equal to.
- 3) (Deadline constraint) All served packets are served before their deadlines. Formally, let $\zeta_l^\mu(n)$ be the time slot in which the n th packet on link l is scheduled by μ if that packet is ever scheduled, and $\zeta_l^\mu(n) = 0$ if that packet is never scheduled by μ . Then the deadline constraint can be stated as follows: For any n and any l with $\zeta_l^\mu(n) > 0$,

$$b_l(n) \leq \zeta_l^\mu(n) < b_l(n) + \tau_l(n).$$

In this paper, we will consider a greedy scheduling policy, called Largest-Deficit-First (LDF) [1] based on the following deficit process $D^\mu(t)$ (also known as debts or virtual queues)

- 1) (Initialization) $D_l^\mu(0) = 0$ for all $l \in \mathcal{K}$.
- 2) (Dynamics) The dynamics of the deficits include the arrival and departure of the deficits.

The arrival of the deficits are based on the arrival of the real packets and a *coin tossing process* that determines whether the arriving packet is counted as a deficit arrival or not. Let the coin tossing process for link l , denoted by $\{C_l(n), n \geq 1\}$, be an i.i.d. Bernoulli process with mean p_l . It is assumed that $\{C_l(n), n \geq 1\}$ is independent across l . Let $B_l(t)$ be the cumulative deficit arrival on link l given by $B_l(0) = 0$ and

$$B_l(t) - B_l(t-1) = \sum_{n=A_l(t-1)+1}^{A_l(t)} C_l(n),$$

where by definition $B_l(t) - B_l(t-1) = 0$ if $A_l(t-1) = A_l(t)$. Then each packet arrival is counted in the cumulative deficit arrival with probability p_l .

The deficit decreases by one each time a packet is scheduled until it reaches zero. Hence, the evolution of the deficit process for link l is then given by

$$D_l^\mu(t) = [D_l^\mu(t-1) + (B_l(t) - B_l(t-1)) - (S_l^\mu(t) - S_l^\mu(t-1))]^+,$$

where $(\cdot)^+ = \max\{0, \cdot\}$.

Observe from the definition that the deficit process keeps track on the amount of service we owe to a link in order to fulfill the minimum delivery rate. To see that, note that the arrival rate of deficit on link l is $\bar{a}_l p_l$. The deficit of link l reduces by one when a packet is successfully transmitted over link l before its deadline. So if all deficits are bounded, then the requirements on packet minimum delivery rates are fulfilled.

The LDF scheduling policy is defined as follows. At each time slot, LDF first sorts the links \mathcal{K} according to the current deficits D with arbitrary tie-breaks, and gets the index vector I such that $D_{I_1} \geq D_{I_2} \geq \dots \geq D_{I_K}$. LDF starts with the selection $\mathcal{E} = \{I_1\}$, which only consists of the link with the largest deficit. Then LDF repeatedly considers the link with the next largest deficit I_i for i from 2 to K and adds it into the selection \mathcal{E} if the following two conditions are satisfied:

- 1) Link I_i does not interfere with any link in \mathcal{E} ; i.e., there exists some $M_r \in M$ such that M_r schedules both \mathcal{E} and I_i .
- 2) There is at least one packet available for transmission on link I_i ; i.e., $Q_{I_i} > 0$, where Q_l is the number of available packets on link l .

The procedure ends when all links have been considered, and the final selection of links is the desired LDF schedule.

III. PRELIMINARIES

In this section, we introduce basic definitions on stability and efficiency ratio that will be used in the following sections. We first define the stability of the system, which was first proposed by Loynes [13] as substability, and was used by, e.g., Maguluri et al. [14] and Srikant and Ying [15] (Chapter 4.2).

Definition 1: The system is *stable* under a scheduling policy μ if the corresponding deficit process $\{D^\mu(t), t \geq 0\}$ satisfies

$$\lim_{C \rightarrow \infty} \limsup_{t \rightarrow \infty} \Pr \left(\sum_{l \in \mathcal{K}} D_l^\mu(t) \geq C \right) = 0.$$

Note that if $\{D(t), t \geq 0\}$ is an aperiodic, irreducible and positive recurrent Markov chain, then the system is stable as defined in Definition 1, since the sum of the deficits converges in distribution as time goes to infinity.

Obviously, the stability of the system depends on the arrival distributions given by $f(\cdot)$, the maximum delay distribution given by $g(\cdot)$, and the required minimum delivery rate $p = (p_l: l \in \mathcal{K})$. Without loss of generality, we fix f and g and consider the stability of the system in terms of the deficit arrival rate $\lambda = (\lambda_l: l \in \mathcal{K})$ with $\lambda_l = \bar{a}_l p_l$. We then have the following definition for characterizing such a relation.

Definition 2: The deficit arrival rate vector λ is *supportable* by a scheduling policy if the system is stable under that policy with deficit rate λ_l for each link l .

Definition 3: The *stability region* of a scheduling policy μ is

$$\Lambda_\mu = \{\lambda \succeq 0: \lambda \text{ is supportable by } \mu\},$$

where \succeq denotes pairwise greater than or equal to.

Note that the stability region defined here is different from the conventional stability region for non-real-time traffic as it is in terms of deficit arrival rates rather than packet arrival rates. This is due to the constraint that packets cannot be scheduled after their deadlines, which makes the stability of the system depend on the specific distributions of packet arrivals and deadlines. As a result, we investigate the stability by fixing f and g while varying the QoS p .

Let the set of all causal scheduling policies be \mathcal{M} , where a causal scheduling policy, also known as an online policy, is one that makes decision based on current information but not future information. We then have the following characterization.

Definition 4: The *maximum stability region* of the system is

$$\Lambda = \bigcup_{\mu \in \mathcal{M}} \Lambda_\mu.$$

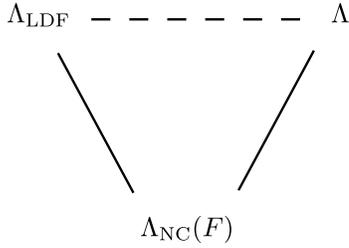


Fig. 2: A roadmap of the proof of Theorem 1.

That is, the maximum stability region of the system is the set of deficit arrival vectors that can be supported by some causal scheduling policy.

For a given scheduling policy μ , the efficiency ratio of the scheduling policy is defined as follows.

Definition 5: The *efficiency ratio* of a scheduling policy μ is

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

While refined characterizations of the stability region are possible [16], [17], the efficiency ratio is still a critical metric to evaluate the throughput performance of a scheduling policy.

IV. MAIN RESULTS

In this section we present the main results of the LDF policy for scheduling real-time traffic in wireless networks. The first result is Theorem 1, which provides a lower bound on the efficiency ratio of the LDF policy, called the real-time local-pooling factor (R-LPF), in the case when the traffic distributions are known. The second result is Theorem 2, which gives lower bounds on the R-LPF regardless of the traffic distributions by assigning weights to the links and calculating the ratio of the weighted sum of the LDF schedule to that of the optimal schedule.

We provide a roadmap of the proof of Theorem 1 in Fig. 2. The goal of Theorem 1 is to establish the connection between Λ , the maximum stability region of the system, and Λ_{LDF} , the stability region of the LDF policy. However, characterizing Λ turns out to be extremely difficult due to the general arrival and maximum delay distributions. We therefore have to introduce a region called $\Lambda_{NC}(F)$, which is the maximum stability region by dividing the time into frames with length F and assume (i) all information within a frame (arrivals and maximum delays) are known at the beginning of a frame and (ii) at the end of a frame, all packets that have not been transmitted are dropped. The region is denoted by $\Lambda_{NC}(F)$ since the frame length is F and the system is non-causal because of condition (i).

This novel frame concept was first introduced by Hou et al. [1] for real-time scheduling in wireless networks and provides an analytical framework for understanding real-time communication in wireless networks. The framework has then been extended to general networks and traffic patterns. In particular, the capacity of the non-causal system and heterogeneous deadlines has been characterized by Jaramillo et al. [6]; i.e., $\Lambda_{NC}(F)$ is known.

We will use $\Lambda_{NC}(F)$ to bridge Λ and Λ_{LDF} . In the theorem, we will first show that

$$\text{int}(\Lambda) \subseteq \liminf_{F \rightarrow \infty} \Lambda_{NC}(F),$$

where $\text{int}(\Lambda)$ is the interior of the set Λ and $\liminf_{F \rightarrow \infty} \Lambda_{NC}(F)$ is the limit set of $\Lambda_{NC}(F)$ as F goes to infinity. After that, we will prove that

$$\sigma^* \text{int}(\Lambda_{NC}(F)) \subseteq \Lambda_{LDF},$$

where σ^* is a constant, called the real-time local-pooling factor whose definition is presented in Section IV-A. Combining the two results together, we will be able to prove that σ^* is a lower bound on the efficiency ratio. *We remark that the second step is non-trivial since we will compare the time-slot-based, causal LDF (not frame-based LDF) with the frame-based, non-causal system.*

As for the second result regarding lower bounds on the R-LPF, we first reformulate the dual problem of solving the R-LPF as a weight assignment problem following Li et al. [16] In the weight assignment problem we try to maximize the ratio of the smallest weighted sum of a schedule to the largest one over different weights within a frame. A similar result for the special case of all-one weight assignment has been observed by Reddy et al. [8] for characterizing the local-pooling factor for fading channels. We then look at the multigraph of the network for any given traffic pattern in a frame, and obtain further lower bounds on the R-LPF using the local weight ratios which do not require the traffic information. Our result immediately implies that the R-LPF is at least $1/(\beta + 1)$ for networks with interference degree β , regardless of the distributions of the packet arrivals and the maximum delays.

A. Real-Time Local-Pooling Factor

We will define a quantity analogous to the local-pooling factor proposed by Joo et al. [12] and the fading local-pooling factor studied by Reddy et al. [8] Before we do that, we need the following two definitions.

Definition 6: A non-causal frame-based scheduling policy μ with frame size F (called an *F-framed* policy for abbreviation) is defined as follows. The packet arrivals and deadlines in the k th frame are known to the policy μ at the beginning of the frame and all packets that arrive during the k th frame are dropped at the end of the frame if not transmitted. Formally, for any $l \in \mathcal{K}$ and positive integer n with $\zeta_l^\mu(n) > 0$, there exists a positive integer k such that

$$kF + 1 \leq b_l(n) \leq \zeta_l^\mu(n) \leq (k + 1)F,$$

where $\zeta_l^\mu(n)$ was defined in Section II in the deadline constraint condition.

Let the set of all F -framed policies be $\mathcal{M}_{NC}(F)$. Note that $\mathcal{M}_{NC}(F)$ is not a subset of \mathcal{M} since policies in $\mathcal{M}_{NC}(F)$ can be non-causal. The frame concept (alternatively called intervals or periods) has been used in the literature for tractable analytical analysis of delay constrained traffic [2]–[6], where packets that arrive in a frame have deadlines in the same frame. In this paper, we adopt this concept to derive the real-time local-pooling factor for the general traffic model.

Definition 7: The maximum stability region of F -framed policies for a positive integer F is

$$\Lambda_{\text{NC}}(F) = \bigcup_{\mu \in \mathcal{M}_{\text{NC}}(F)} \Lambda_{\mu}.$$

We now introduce some notations needed for the main results. Let $\mathcal{J}(F)$ be the set of arrival and maximum delay patterns in a frame of F time slots. We will call an element of $\mathcal{J}(F)$ an F -pattern. An F -pattern is represented by $J = (A^{(F)}, \tau^{(F)})$ with $A^{(F)} = (A_l^{(F)}(t) : l \in \mathcal{K}, 1 \leq t \leq F)$ and $\tau^{(F)} = (\tau_l^{(F)}(n) : l \in \mathcal{K}, 1 \leq n \leq A_l^{(F)}(F))$, where $A_l^{(F)}(t)$ is the cumulative packet arrival to link l by time slot t in the frame, and $\tau_l^{(F)}(n)$ is the maximum delay associated with the n th packet on link l . Thus, $A_l^{(F)}(F)$ is the total number of arrivals in the frame on link l . Due to the i.i.d. distributions of the arrival and maximum delay given by f and g , there is a stationary distribution of the set of F -patterns, denoted by $\pi : \mathcal{J}(F) \rightarrow \mathbb{R}$.

For a given F -pattern $J = (A^{(F)}, \tau^{(F)})$, a *schedule* $s = (s_l(n) : l \in \mathcal{K}, 1 \leq n \leq A_l^{(F)}(F))$ specifies the time slot at which each packet is scheduled to be transmitted (if it ever gets scheduled), where $s_l(n)$ is a nonnegative integer that indicates the n th packet on link l is scheduled at time slot $s_l(n)$ if $s_l(n) \in \{1, 2, \dots, F\}$, and is never scheduled if $s_l(n) = 0$. A schedule s is *feasible* for the F -pattern J if 1) each scheduled packet is scheduled within its feasible scheduling window, 2) at most one packet is scheduled on each link in one time slot, and 3) the set of links with packets scheduled in each time slot forms a feasible link schedule; i.e., it is a subset of $\{l \in \mathcal{K} : M_{l,r} = 1\}$ for some r . Note that the schedule s here is different from the link schedules defined in Section II in that s specifies the scheduling of each packet in the whole frame, and that s needs to take into account the traffic so that no scheduling is allowed before the arrival or after the deadline of a packet. We also say that a schedule s is *maximal* for J if no more packets can be further scheduled (i.e., no $s_l(t)$ can be changed from 0 to a positive integer) without breaking feasibility. We denote the maximal feasible schedules for J by $S^*(J)$.

We define the *total service vector of schedule* s to be the column vector $W(s) = (W_i(s) : i \in \mathcal{K})$ with $W_i(s) = \sum_{n=1}^{A_i^{(F)}} \mathbb{1}_{\{s_i(n) \neq 0\}}$, where $\mathbb{1}_{\{s_i(n) \neq 0\}}$ is the indicator function. Then $W(s)$ is the vector of total number of scheduled packets on each link for the schedule s . Let the *maximal service matrix for* J be

$$M_J = \{W(s) : s \in S^*(J)\},$$

where again M_J represents both the set and the matrix consisting of the vectors as its columns, by abuse of notation. Then the columns of M_J are the total service vectors of the maximal schedules. We note that M_J does not contain all-zero columns if and only if J includes at least one packet arrival on some link, since schedules in $S^*(J)$ are maximal. Similarly, define $M_{J,L}$ to be the maximal service matrix restricted to the set of links L for given pattern J . Then $M_{J,L}$ has no all-zero columns if and only if the pattern J includes at least one packet on some link in L . Also note that $M_{J,L}$ has $|L|$ rows while M_J has K rows.

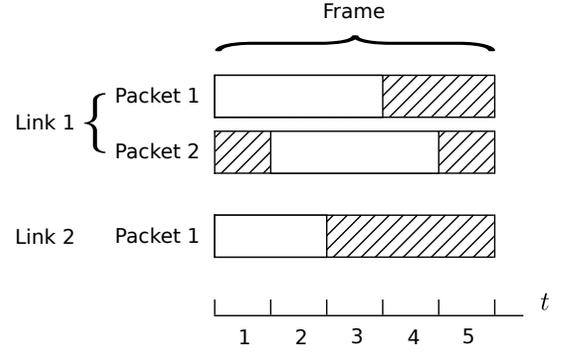


Fig. 3: An example of a 5-pattern for two links.

We use the example in Fig. 3 to illustrate the above notations and the concept of the maximal service matrix. As shown in the figure, we consider a frame with size 5 and a 5-pattern J with two packets arriving to link 1 and one packet arriving to link 2, whose arriving times and deadlines are indicated by the blank bars in the figure. The corresponding pattern can be represented by $J = (A^{(5)}, \tau^{(5)})$, where $\{A_1^{(5)}(t), t \geq 0\} = (0, 1, 2, 2, 2, 2)$, $\{A_2^{(5)}(t), t \geq 0\} = (0, 1, 1, 1, 1, 1)$, $\{\tau_1^{(5)}(n), n \geq 1\} = (3, 3)$, and $\{\tau_2^{(5)}(n), n \geq 1\} = (2)$. Assume the two links interfere with each other, so at each time slot only one of them can be scheduled. We can check that there are eight maximal feasible schedules in $S^*(J)$ as follows:

$$s^1 = \begin{pmatrix} (1, 2) \\ (0) \end{pmatrix}, s^2 = \begin{pmatrix} (1, 3) \\ (2) \end{pmatrix}, s^3 = \begin{pmatrix} (1, 4) \\ (2) \end{pmatrix}, s^4 = \begin{pmatrix} (2, 3) \\ (1) \end{pmatrix},$$

$$s^5 = \begin{pmatrix} (2, 4) \\ (1) \end{pmatrix}, s^6 = \begin{pmatrix} (3, 2) \\ (1) \end{pmatrix}, s^7 = \begin{pmatrix} (3, 4) \\ (1) \end{pmatrix}, s^8 = \begin{pmatrix} (3, 4) \\ (2) \end{pmatrix},$$

where the first row of s^i is the schedule for the two packets on link 1, and the second row is the schedule for the packet on link 2. Then the total service vectors are

$$W(s^1) = \begin{pmatrix} 2 \\ 0 \end{pmatrix} \text{ and } W(s^i) = \begin{pmatrix} 2 \\ 1 \end{pmatrix} \text{ for } 2 \leq i \leq 8.$$

Hence the maximal service matrix is

$$M_J = \begin{pmatrix} 2 & 2 \\ 0 & 1 \end{pmatrix}.$$

We remark from the above example that unlike in the scenario of non-real-time traffic [11], the total service vector of one maximal schedule could be dominated by that of another in the real-time setting. Thus the maximal service matrix M_J can be huge and hard to compute, especially for large frame size F and complex traffic pattern J .

Definition 8: The *real-time local-pooling factor (R-LPF)* for the F -framed scheduling policies for the set of links $L \subseteq \mathcal{K}$ is

$$\sigma_L^*(F) = \inf\{\sigma : \exists \phi_1, \phi_2 \in \Phi_L(F) \text{ such that } \sigma \phi_1 \succeq \phi_2\},$$

where $\Phi_L(F)$ is the *service region restricted to the set of links* $L \subseteq \mathcal{K}$ defined by

$$\Phi_L(F) = \left\{ \phi : \phi = \sum_{J \in \mathcal{J}(F)} \pi(J) \eta_J, \eta_J \in \mathcal{CH}(M_{J,L}) \right\},$$

and $\mathcal{CH}(M_{J,L})$ defines the convex hull over the columns of the matrix $M_{J,L}$.

Definition 9: The R-LPF for the F -framed scheduling policies is

$$\sigma^*(F) = \min_{L \subseteq \mathcal{K}} \sigma_L^*(F).$$

Definition 10: The R-LPF for the system is

$$\sigma^* = \liminf_{F \rightarrow \infty} \sigma^*(F).$$

We then have the following theorem stating that the R-LPF is a lower bound on the efficiency ratio of LDF.

Theorem 1: $\gamma_{\text{LDF}}^* \geq \sigma^*$.

Intuitively when the frame length goes to infinity the loss at the edge of frames becomes negligible. The proof of Theorem 1 uses the strictly separating hyperplane theorem [18] and follows the fluid limit technique that was first proposed by Dai [9] for multiclass queueing systems and later developed for discrete-time generalized switches by Andrews et al. [19] and further used in wireless networks by Reddy et al. [8] and Ji et al. [20] The complete proof is presented in the supplementary materials.

By the definition of the R-LPF, we can get the R-LPF by solving the following linear program for each $L \subseteq \mathcal{K}$, as suggested by Li et al. [16] and Reddy et al. [8]:

$$\begin{aligned} \sigma_L^*(F) = \min_{\sigma, \rho, \theta} \quad & \sigma \\ \text{s.t.} \quad & M_L(F)\theta - M_L(F)\rho \succeq 0 \\ & \mathbf{1}^\top \theta - \sigma = 0 \\ & \mathbf{1}^\top \rho - 1 = 0 \\ & \rho, \theta \in \mathbb{R}_+^r, \end{aligned} \quad (1)$$

where $M_L(F)$ is the vertices of the polygon $\sum_{J \in \mathcal{J}(F)} \pi(J) \mathcal{CH}(M_{J,L})$ (the summation is in the sense of the Minkowski sum and $\mathcal{CH}(M_{J,L})$ denotes the convex hull of the column vectors in $M_{J,L}$, where $r_{J,L}$ denotes the number of columns), $r = \sum_{J \in \mathcal{J}(F)} r_{J,L}$ is the cardinality of $M_L(F)$, ρ and θ are nonnegative column vectors of length r , $\mathbf{1}$ is the all-one column vector of length r , $(\cdot)^\top$ denotes the transpose, and \mathbb{R}_+ denotes the set of nonnegative real numbers. That said, computing the exact R-LPF is usually complex, as it involves roughly

$$\sum_{L \subseteq \mathcal{K}} \left(2 \sum_{J \in \mathcal{J}(F)} r_{J,L} + |L| \right)$$

constraints for each F , which increases exponentially with both the size of the network K and the frame size F . Thus, we seek lower bounds on the R-LPF in the next subsection.

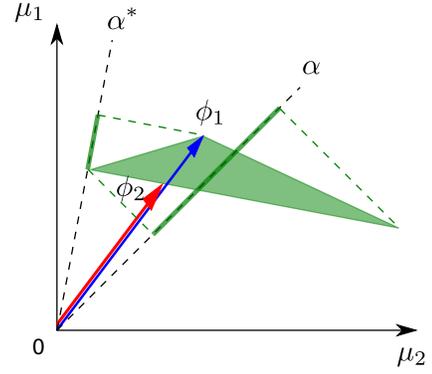


Fig. 4: Illustration of the projection interpretation of the dual formulation of $\omega_L^*(F)$ with $L = \mathcal{K} = \{1, 2\}$. The shaded triangular area is the convex hull of $M_L(F)$.

B. Characterizing the R-LPF

1) *Dual of the R-LPF:* The dual problem of (1) is given by the following [16]

$$\begin{aligned} \omega_L^*(F) &= \max_{\alpha_L, \omega} \quad \omega \\ \text{s.t.} \quad & \mathbf{1}^\top \succeq \alpha_L^\top M_L(F) \succeq \omega \mathbf{1}^\top \\ & \alpha_L \in \mathbb{R}_+^{|L|} \\ & = \max_{\alpha_L \in \mathbb{R}_+^{|L|}} \frac{\min_{\phi \in M_L(F)} \alpha_L^\top \phi}{\max_{\phi \in M_L(F)} \alpha_L^\top \phi}, \end{aligned} \quad (2)$$

where we adopt the useful convention $\frac{0}{0} = 0$. By the strong duality of the linear program (1), finding the R-LPF for subset L and frame size F is equivalent to the weight assignment problem (2); i.e., $\sigma_L^*(F) = \omega_L^*(F)$.

Another interpretation of the dual problem is that since $\alpha_L^\top \phi$ is the length of the projection of ϕ along the vector α_L (module the length of α_L), the optimization problem (2) is to find the best projection direction α_L for each subset L such that the ratio of the smallest projection to the largest one from $M_L(F)$ is maximized. Note that for any α_L we have a corresponding ratio of the smallest to the largest projections, which is a valid lower bound on $\sigma_L^*(F)$. We illustrate this interpretation in Fig. 4 with $L = \mathcal{K} = \{1, 2\}$. In the figure, α^* is the optimal projection direction since the ratio of the smallest vector (ϕ_2 projected to the direction of α^*) to the largest vector (ϕ_1 projected to the direction of α^*) in the projection along the direction of α^* is the maximum (it equals $\omega_L^*(F)$) among all the possible projection directions. The vector α in the figure is an arbitrary direction. The ratio of the smallest vector projection to the largest vector projection is smaller than that along the direction α^* , as shown in the shaded segment along α , and thus provides a lower bound on $\omega_L^*(F)$.

2) *Lower Bounds for Conflict Graph Interference Model:* While the dual problem (2) gives a different view of the original problem for solving the R-LPF, the problem is not simplified since the size of $M_L(F)$, where $M_L(F)$ is taken as the set of columns, grows exponentially with F . As a result, we are interested in finding a lower bound on the R-LPF that can be computed efficiently (without calculating the Minkowski sum of the maximal service matrices for all traffic

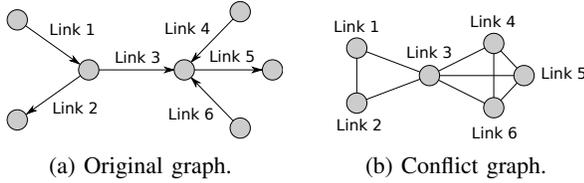


Fig. 5: An original graph and its conflict graph.

pattern and all frame size). In this subsection we introduce the lower bounds on R-LPF for networks with interference models that are represented by conflict graphs [21] (also known as interference graph), where two links either interfere with each other exclusively, or do not interfere at all. An example of the original graph and its conflict graph is given in Fig. 5. We introduce the ideas of pressure and minimum pressure in the following, which use the local information to estimate the global optimal value in (2).

For any $\alpha \in \mathbb{R}_+^K$, we define the *pressure* of link i to be

$$\kappa_i(\alpha) = \frac{\alpha_i}{\alpha_i + \max_{I \in \mathcal{I}(i)} \sum_{j \in I} \alpha_j}, \quad (3)$$

where $\mathcal{I}(i) \subseteq \mathcal{P}(N(i))$ is the collection of subsets of the neighbors of i that can be scheduled simultaneously ($N(i)$ is the set of links that interfere with link i , and $\mathcal{P}(\cdot)$ denotes the power set). We also define the *minimum pressure* for given vector $\alpha \in \mathbb{R}_+^K$ by

$$\psi(\alpha) = \min_{i \in \mathcal{K}} \kappa_i(\alpha). \quad (4)$$

So $\psi(\alpha)$ is just the lowest pressure for α over all links. Then we have the following lower bound on the R-LPF.

Theorem 2:

$$\sigma^* \geq \sup_{\alpha \in \mathbb{R}_+^K} \psi(\alpha).$$

Intuition and proof outline. Given an arbitrary vector of nonnegative weights α on the links, we define

$$G(\alpha, L, F) = \frac{\min_{\phi \in M_L(F)} \alpha_L^T \phi}{\max_{\phi \in M_L(F)} \alpha_L^T \phi},$$

where α_L is the vector of α restricted to the subset L . Then $G(\alpha, L, F)$ is the minimum global weight ratio of schedules for subset L and frame size F . By the dual representation in (2), the R-LPF is lower-bounded by the smallest $G(\alpha, L, F)$ over all possible L and F . Note that while $G(\alpha, L, F)$ is defined over $M_L(F)$ which is averaged over all possible traffic patterns in $\mathcal{J}(F)$ according to the traffic distributions, we can work on an arbitrary traffic $J \in \mathcal{J}(F)$ to establish a universal lower bound that holds for any $J \in \mathcal{J}(F)$, which will also be a lower bound on $G(\alpha, L, F)$. Since

$$G(\alpha, L, F) \geq \frac{\min_{\phi \in M_{J,L}} \alpha_L^T \phi}{\max_{\phi \in M_{J,L}} \alpha_L^T \phi}$$

for any $J \in \mathcal{J}(F)$, we only need to lower-bound the ratio of the weights of two maximal schedules given the specific traffic. Using a *multigraph* representation of the network, for

any $J \in \mathcal{J}(F)$ and any $\phi_1, \phi_2 \in M_{J,L}$, we can divide the scheduled packets of ϕ_1 and ϕ_2 into groups such that

$$\frac{\alpha_L^T \phi_1}{\alpha_L^T \phi_2} \geq \frac{\sum_{i=1}^N x_i}{\sum_{i=1}^N y_i},$$

where N is the total number of packets scheduled in ϕ_1 , x_i is the total weight of the i th packet in ϕ_1 , and y_i is the weight of neighboring packets of the i th packet of ϕ_1 scheduled by ϕ_2 . By the group construction and the definition of minimum pressure $\psi(\alpha)$, we have $\frac{x_i}{y_i} \geq \psi(\alpha)$ for any i regardless of L and F . Then $G(\alpha, L, F)$ is lower-bounded by the minimum pressure $\psi(\alpha)$. Since the minimum pressure $\psi(\alpha)$ is determined by the network topology and does not depend on L or F , we get a lower bound on the R-LPF, which can then be strengthened to the form of Theorem 2 by optimizing over α . The detailed proofs can be found in Appendix A.

Remark 1. We emphasize that while the R-LPF involves computing the maximal service matrices for all subsets of \mathcal{K} and all traffic patterns under all frame sizes, Theorem 2 states that the maximum value of the minimum pressure, which is a purely topological quantity for the entire network, serves as a lower bound on the R-LPF.

Remark 2. Theorem 2 implies that any $\alpha \in \mathbb{R}_+^K$ will give a lower bound on the R-LPF by the corresponding minimum pressure $\psi(\alpha)$. In particular, we have the following corollary.

Corollary 1: For a network with an interference degree of β ,

$$\sigma^* \geq \frac{1}{\beta + 1}.$$

Corollary 1 follows directly from Theorem 2 by setting $\alpha = \mathbf{1}$.

We note that this translates to $\sigma^* \geq 1/3$ in the one-hop interference model, where the interference degree is at most 2 (this can be easily proved by noticing that the neighboring links of link 3 in Fig 5 form two cliques in the conflict graph). This is related to the well-known result that LQF has efficiency ratio at least 1/2 in packet switches [22], [23] and in wireless networks under the one-hop interference [24], [25], where $\beta = 2$. The lower bound on the efficiency ratio of the greedy scheduling policy decreases from 1/2 in the non-real-time case to 1/3 in the real-time case due to the temporal correlation among packets brought by deadlines, which does not exist in non-real-time traffic. This can be illustrated by considering a star network with one center link and two leaf links, where the center link interferes with any leaf link but the two leaf links do not interfere each other. Suppose one packet arrives at the center link at the beginning of time slot 1 with deadline at the end of time slot 2, and one packet arrives at each of the two leaf links at the beginning of time slot 1 with deadline at the end of time slot 1. Then the optimal scheduler will schedule the two leaf links at time slot 1 and the center link at time slot 2, while LDF may schedule the center link at time slot 1 and nothing at time slot 2 (since the packets on the leaf links have already expired), which results in an efficiency ratio of 1/3. Note that if the above arrival pattern is for non-real-time packets (i.e., there is no deadline of the packets), then the longest-queue-first (LQF) can schedule at least one

packet at each time slot when at least one of the queues is not empty, while the optimal scheduler may schedule at most two packets at each time slot, so LQF guarantees at least half throughput. To sum up, LDF has a smaller efficiency ratio lower bound than its non-real-time counterpart LQF because it may inevitably schedule the “wrong” packets due to its inability to take into account the consequences in the future of its current decisions, and this cannot be compensated by future actions.

3) *Lower Bounds on R-LPF for Special Networks:* We now do a case study of lower bounds on R-LPF for some special networks via the minimum pressure technique stated in Theorem 2.

a) *Collocated Network:* We first consider the scenario of the collocated network, where at most one link can be scheduled at each time slot. Notice that the interference degree of the network is $\beta = 1$ since in any subset of the links there is at most one link that can be scheduled. Then by Theorem 1 and Theorem 2, the efficiency ratio is at least $1/2$.

b) *Star Networks:* Consider a star network with interference degree β . Then this network consists of one center link and β leaf links, where at each time slot either the center link or all the leaves can be scheduled. By setting the weight of the center link to be $\sqrt{\beta}$ and the weight of each leaf to be 1, we get the minimum pressure $\psi = 1/(\sqrt{\beta} + 1)$. Hence the efficiency ratio is at least $1/(\sqrt{\beta} + 1)$ for the star network with interference degree β .

c) *Tree Networks:* For tree networks with interference degree β , a lower bound on the R-LPF is given in the following corollary.

$$\text{Corollary 2: } \sigma^* \geq \frac{1}{2\sqrt{\beta-1}+1}.$$

The proof can be found in Appendix C.

Remark. Note that the interference degree is equal to the largest degree of a link in the trees. Also note that this lower bound is better than the $1/(\beta + 1)$ minimum pressure bound given by the all-one vector for $\beta \geq 3$.

V. THE CONSENSUS ALGORITHM

We design an algorithm that can be used to compute a lower bound on the R-LPF. Let each link i maintain a weight α_i and a pressure κ_i . In each time slot all links broadcast (α_i, κ_i) to its neighbors and update its weight and pressure by

$$\Delta\alpha_i = z \sum_{j \in N(i)} (\kappa_j - \kappa_i)$$

and

$$\kappa_i = \frac{\alpha_i}{\alpha_i + \max_{I \in \mathcal{I}(i)} \sum_{j \in I} \alpha_j}.$$

The constant z can be interpreted as the step size. The intuition behind the simple heuristic algorithm is that if under current weight assignment one link has pressure greater than those of its neighbors, then the weight of that link should be transferred to its neighbors so that the minimum of their pressures can increase. Likewise, when one link has pressure less than those of its neighbors, then the weights on its neighbors should be transferred to that link to make the minimum pressure higher. We would expect the algorithm to converge as time goes

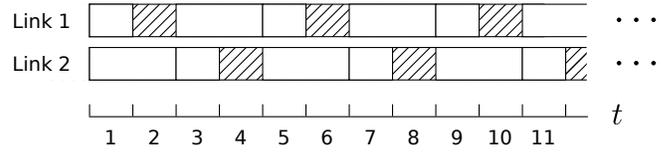


Fig. 6: An adversarial traffic pattern for a collocated network with two links. Each blank bar indicates the arriving time and deadline of a real-time packet. The packets arrive on link 1 at the beginning of time slots 1, 3, 5, 7, ..., and must be scheduled before the end of time slots 1, 4, 5, 8, The packets arrive on link 2 at the beginning of time slots 1, 3, 5, 7, ..., and must be scheduled before the end of time slots 2, 3, 6, 7,

by when the step size is sufficiently small; i.e., the weights and pressures for all links remain unchanged eventually. Then the weight vector that our algorithm converges to yields a minimum pressure that lower-bounds the R-LPF. Note that we call this algorithm “consensus” because the pressures for the links will usually converge to the same value and reach consensus. We evaluate the performance of the consensus algorithm in the Section VII.

VI. DISCUSSIONS

A. Efficiency Ratios Under Adversarial Traffic

In this section we discuss the performance of LDF under adversarial traffic. We consider a more general type of traffic, where, instead of i.i.d., the packet arrival, maximum delay, deficit arrival and tie-breaking processes are relaxed to be irreducible positive recurrent Markov chains as by Andrews et al. [19] We illustrate that when the traffic is adversarial in this general type, the efficiency ratio of LDF can be as low as $1/(\sqrt{\beta} + 1)$, where β is the interference degree. In particular, for collocated networks the efficiency ratio of LDF under adversarial traffic is consistent with the lower bound given in Corollary 1.

We start with a collocated network with two links. Consider the adversarial traffic given in Fig. 6. Assume that the deficits for both links are the same at the beginning of time slot 1. Also assume that when there is a packet arriving to each link (time slots 1, 3, 5, 7, ...), the deficits on both links increase by one with probability $1/2 + \epsilon$ for some small positive ϵ , and remain unchanged with probability $1/2 - \epsilon$. This results in minimum delivery rates $p_i = 1/2 + \epsilon$ for $i = 1, 2$. We further assume that when the deficits on the two links are equal, the tie-breaking rule of LDF gives priority to link 2. Then one can easily see that LDF schedules link 2 at time slots 1, 5, 9, ..., and schedules link 1 at time slots 3, 7, 11, ..., while LDF idles at even time slots. Then the average deficit arrival to each link per time slot is $1/4 + \epsilon/2$, and the average deficit departure from each link per time slot is $1/4$. Hence the deficits are not stable under LDF given this traffic pattern. However, one would notice that the optimal scheduler could schedule link 1 in time slots $4k$ and $4k + 1$ and schedule link 2 in time slots $4k + 2$ and $4k + 3$, for all positive integer k . Hence the optimal scheduler can stabilize the system when the minimum

delivery rates are $p_i = 1$ for $i = 1, 2$. By making ϵ arbitrarily small we can see that the efficiency ratio of LDF is at most $1/2$ in this two-link collocated network, which meets with the lower bound given in Corollary 1.

We now consider a general network. In the following theorem we construct an adversarial traffic process.

Theorem 3: There exists a traffic pattern distribution such that $\gamma_{\text{LDF}}^* \leq \frac{1}{\sqrt{\beta+1}}$, where β is the interference degree.

The theorem can be proved by finding the link with interference degree β and constructing a specific traffic pattern on that graph. The detailed proof can be found in Appendix B.

B. Extension to Heterogeneous Link Rates

The LDF policy can be generalized to heterogeneous integer-valued link rate scenario following Dimakis and Walrand [11]. Assume the link rate for link i is $c_i \in \mathbb{N}_+$ for $i \in \mathcal{K}$. LDF now schedules c_i packets instead of 1 packet on each selected link i . Then Theorem 1 still holds by replacing the 1's on the i th row of M with c_i , and Theorem 2 still holds after redefining the pressure by

$$\kappa_i(\alpha) = \frac{\alpha_i}{\alpha_i + \max_{I \in \mathcal{I}(i)} \sum_{j \in I} \alpha_j c_j}.$$

Note that in the summation of the denominator all the weights α_j 's are multiplied by the corresponding link rate c_j , while α_i in both the denominator and the numerator are not multiplied by c_i . Intuitively this is due to the fact that in the worst case LDF can schedule only one packet on link i as opposed to $\max_{I \in \mathcal{I}(i)} \sum_{j \in I} c_j$ packets scheduled by the optimal policy on the neighboring links of i . The consensus algorithm can also be modified according to the new definition of pressure.

VII. SIMULATIONS

In this section we use simulations to evaluate the stability performance of LDF, as well as the consensus algorithm we proposed.

A. Stability Performance

Since to the best of our knowledge, neither the maximum stability region nor an optimal scheduling policy has been obtained in the literature, we do not have a benchmark for the stability performance of LDF. As a result, we compare LDF to two other scheduling policies that do not depend on frames and evaluate the performance using simulations. The first simple scheduling policy we consider is RandMax, which randomly chooses a maximal schedule over the links with packets in each time slot. The other one is MaxWeight, which chooses a maximal schedule with the maximum deficit sum over the links with packets in each time slot.

We first considered a 4-link linear network with one-hop interference. We assumed the packet arrival distribution is binomial with number of trials 2 and success probability 0.5, and the maximum delay distribution is uniform over $\{2, 3, 4\}$. This gives us packet arrival rate $\bar{a} = 1$ and mean maximum delay $\bar{\tau} = 3$. We varied the minimum delivery rate to vary the deficit arrival rate. We compare the average deficit sums of the last 1,000 iterations under the three policies, where

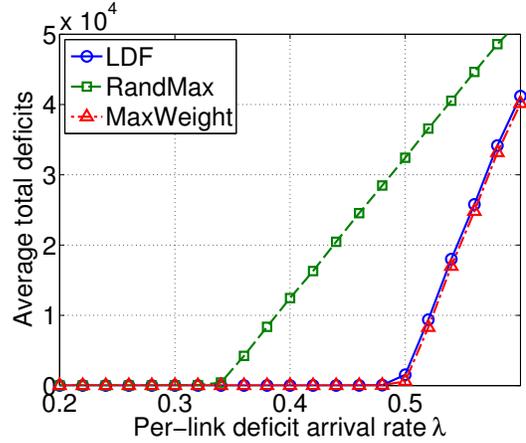


Fig. 7: Comparison of the three scheduling policies on a four-linear network with one-hop interference.

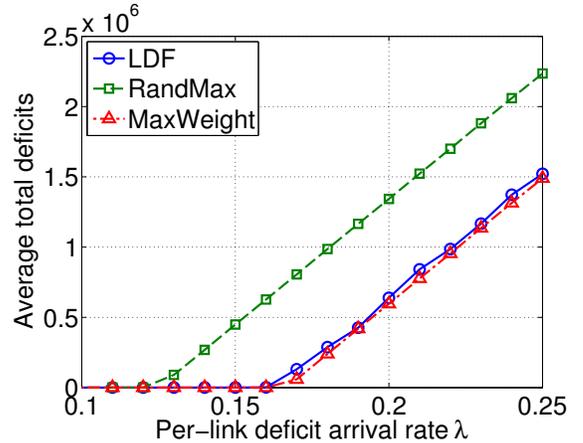


Fig. 8: Comparison of the three scheduling policies on a nine-cycle network with two-hop interference.

each simulation is run for 100,000 iterations. The results are shown in Fig. 7. As can be observed from the figure, LDF and MaxWeight have similar stability performance, achieving a maximum deficit arrival rate of roughly 0.5 and significantly outperform the simple RandMax policy, which achieves a maximum deficit arrival rate of roughly 0.33. We further remark that for non-real-time traffic, the maximum deficit arrival rate is 0.5. Thus both LDF and MaxWeight have a near-optimal performance in this case.

We also consider a nine-cycle network with two-hop interference, whose non-real-time local-pooling factor is $2/3$. The arrival and deadline distributions are the same as the previous case, and the number of iterations is 100,000. The results are shown in Fig. 8. Note that in this example, RandMax is still the worst of the three, achieving a maximum deficit arrival rate roughly 0.12, while MaxWeight is slightly better than LDF, both of which achieve a maximum deficit arrival rate roughly 0.16. We note that for non-real-time traffic the maximum deficit arrival rate is $1/3$. As we have been trying to convey in this paper, the maximum stability region for the specific packet arrival and deadline distribution is unknown. We only know

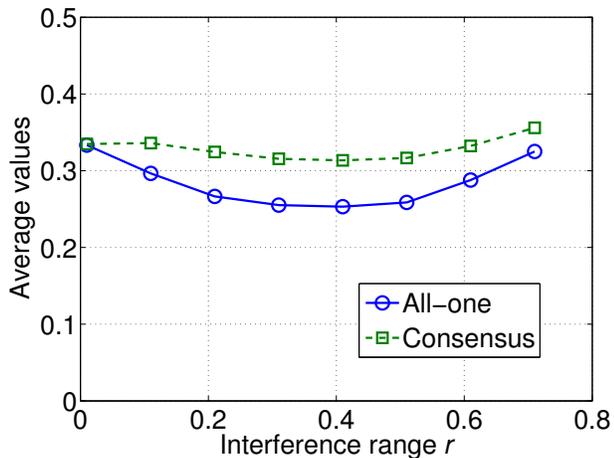


Fig. 9: Lower bounds on the R-LPF given by the consensus algorithm and the all-one algorithm for different interference range.

that the maximum rate for the real-time traffic is $\bar{\lambda} \leq 1/3$. Note that the nine-cycle has an interference degree of 2, so by Theorem 2, LDF has an efficiency ratio of at least $1/3$, which agrees with the simulation result since $0.16 > \frac{1}{3} \times \frac{1}{3} \geq \frac{1}{3}\bar{\lambda}$.

Therefore, both simulations imply good throughput performance of LDF and validate our lower bound on the efficiency ratio.

B. Performance of the Consensus Algorithm

We now study the performance of the consensus algorithm we proposed in Section V. We run the consensus algorithm on random networks with the unit-disk interference model used by Joo et al. [12]

We place 32 nodes randomly in a unit square area. Any two nodes with distance less than the communication range $r_c = 0.25$ may form a link. The default maximal number of links is 24 (uniformly chosen from possible links). Any two links with minimum node distance less than the interference range r_i interfere each other, and the default interference range is $r_i = 0.4$.

We run the consensus algorithm for 3000 iterations with step size $z = 0.1$ for each network. We compare the average lower bound obtained via the consensus algorithm to the all-one algorithm in Fig. 9. Each point is the average of 100 random networks. We also attach one example network for each interference range r_i in Fig. 10. We see from Fig. 9 that our consensus algorithm achieves a much better lower bound than the one given by the interference degree alone.

We show an example of traces of the consensus algorithm for a random unit-disk network with $r_c = 0.25$ and $r_i = 0.11$ in Fig. 11. The lines in the upper figure of Fig. 11 are the weights α_i 's for the links $i \in \mathcal{K}$ with respect to the iterations, and the lower figure of Fig. 11 shows the minimum pressure given by the weights with respect to the iterations. We see that the weights become unchanged after about 300 iterations, which indicates that the consensus algorithm converges reasonably fast. We also note that the minimum pressure given

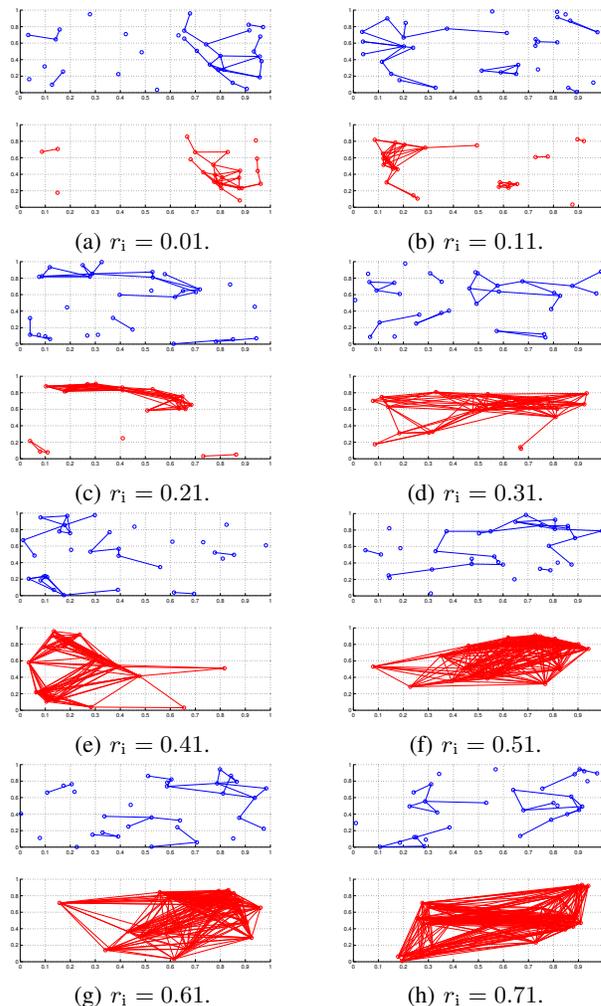


Fig. 10: Example networks for different interference range r_i . The top figures are the original node-link graphs, while the bottom ones are the corresponding conflict graphs.

by the consensus algorithm after 500 is about 0.3, while by Corollary 1 the all-one vector gives a minimum pressure of 0.25 since the interference degree of the tested network is 3.

VIII. CONCLUSIONS

In this paper we considered the problem of scheduling real-time traffic in wireless networks under general stochastic arrivals and deadlines and general interference model. The fraction of delivered packets at a link is required to be no less than a certain threshold. We used deficits to inspect the stability of the system, and studied the stability performance of a scheduling policy that we call the largest-deficit-first (LDF) policy. We proved that the efficiency ratio of LDF can be lower bounded by a quantity that we call the real-time local-pooling factor (R-LPF). Furthermore, we showed lower bounds on the R-LPF can be calculated by assigning weights to the links, with a special case lower bound of $1/(\beta + 1)$, where β is the interference degree. We also proposed a heuristic consensus algorithm that can be used to estimate the R-LPF for general networks.

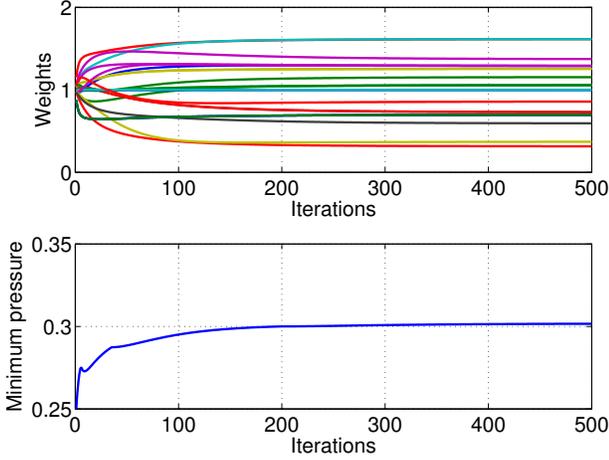


Fig. 11: Traces of weights and minimum pressure under the consensus algorithm.

APPENDIX A PROOF OF THEOREM 2

For given $\alpha \in \mathbb{R}_+^K$, let α_L denote the vector α restricted to the subset $L \subseteq \mathcal{K}$. We define the pressure of link i in the subset $L \subseteq \mathcal{K}$ for vector α to be

$$\kappa_{i,L}(\alpha) = \frac{\alpha_i}{\alpha_i + \max_{I \in \mathcal{I}(i,L)} \sum_{j \in I} \alpha_j},$$

where $\mathcal{I}(i,L) \subseteq \mathcal{P}(N_L(i))$ is the collection of subsets of the neighbors of i in L that can be scheduled simultaneously. We also define the minimum pressure in the subset $L \subseteq \mathcal{K}$ for α to be

$$\psi_L(\alpha) = \min_{i \in L} \kappa_{i,L}(\alpha).$$

We can check with (3) and (4) that

$$\kappa_{i,\mathcal{K}}(\alpha) = \kappa_i(\alpha), \psi_{\mathcal{K}}(\alpha) = \psi(\alpha).$$

We note the following lemma.

Lemma 1: For any $\alpha \in \mathbb{R}_+^K$ and any $L \subseteq \mathcal{K}$,

$$\psi_L(\alpha) \geq \psi(\alpha).$$

Proof:

$$\begin{aligned} \psi_L(\alpha) &= \min_{i \in L} \kappa_{i,L}(\alpha) \\ &= \kappa_{i^*,L}(\alpha) \\ &= \frac{\alpha_{i^*}}{\alpha_{i^*} + \sum_{j \in I^*} \alpha_j} \\ &\geq \frac{\alpha_{i^*}}{\alpha_{i^*} + \sum_{j \in I^{**}} \alpha_j} \\ &\geq \min_{i \in \mathcal{K}} \kappa_i(\alpha) \\ &= \psi(\alpha), \end{aligned}$$

for some $i^* \in L$, some $I^* \in \mathcal{I}(i^*,L)$ and some $I^{**} \in \mathcal{I}(i^*,\mathcal{K})$. \blacksquare

The following lemma is the key to the proof of Theorem 2.

Lemma 2: For any $\alpha \in \mathbb{R}_+^K$, any F , any $J \in \mathcal{J}(F)$ and any $L \subseteq \mathcal{K}$,

$$\frac{\min_{\phi \in M_{J,L}} \alpha_L^T \phi}{\max_{\phi \in M_{J,L}} \alpha_L^T \phi} \geq \psi(\alpha).$$

Proof of Lemma 2: We fix $J \in \mathcal{J}(F)$ and $L \subseteq \mathcal{K}$ and focus on the arrival and maximum delay pattern given by J restricted to the subset of links L . If $\psi(\alpha) = 0$ then the result is trivial, so we may assume $\psi(\alpha) > 0$; i.e., $\alpha \succ 0$. For each link $l \in L$, replace it with n links (each of which has a single packet arrival in the frame) if the total number of packets arriving on l in the frame is $n \geq 2$, leave it alone if the total number of packets arriving on l in the frame is 1, and remove it from our consideration if no packet arrives in this frame according to J . We then get a multigraph whose set of links is denoted by \mathcal{K}' , where $K' = |\mathcal{K}'|$ equals the total number of packets arriving on L in the original conflict graph according to J , and each link in \mathcal{K}' represents a packet in the original conflict graph with arriving time and deadline given by J . The interference model of \mathcal{K}' inherits from the interference model of \mathcal{K} , plus that two links in \mathcal{K}' that correspond to the same link in \mathcal{K} interfere each other. Let $I'(l)$ denote the set of links that interfere with link l in \mathcal{K}' , and by convention assume $l \in I'(l)$. Also let $\mathcal{I}'(l)$ be the collection of subsets of $I'(l) \setminus \{l\}$ that can be scheduled simultaneously according to the interference model.

A multi-schedule over the multigraph \mathcal{K}' in the frame is represented by a function (we overload the symbol s for convenience in this proof)

$$\begin{aligned} s: \mathcal{K}' \times \{1, 2, \dots, F\} &\rightarrow \{0, 1\} \\ \langle i, t \rangle &\mapsto s_i(t) \end{aligned}$$

with $s_i(t) = 1$ if link $i \in \mathcal{K}'$ is scheduled by s at time slot t , and $s_i(t) = 0$ otherwise. A multi-schedule s is feasible if no two interfering links are scheduled at the same time slot, no link is scheduled before its arriving time or after its deadline, and each link is scheduled at most once during the entire frame. A feasible multi-schedule s is maximal if no more links can be scheduled without breaking the feasibility. We note that a feasible (or maximal, respectively) schedule for J over the original set of links \mathcal{K} corresponds to a feasible (or maximal, respectively) multi-schedule over the multigraph \mathcal{K}' given by J . Let $\text{supp}(s)$ be the support of s , i.e., the set of $\langle \text{link}, \text{time slot} \rangle$ pairs of scheduled links by s . Let $\|s\| = \sum_i (\bar{\alpha}_i \sum_t s_i(t))$, where $\bar{\alpha}_i = \alpha_j$ if link i in \mathcal{K}' corresponds to link j in \mathcal{K} . Then we say $\|s\|$ is the weight of the multi-schedule s .

Define the interference neighborhood of the $\langle \text{link}, \text{time slot} \rangle$ pair $\langle i, t \rangle$ to be the interfering links of link i at time slot t ; i.e.,

$$I'(i, t) = I'(i) \times \{t\} \subseteq X,$$

where $X = \mathcal{K}' \times \{1, 2, \dots, F\}$.

We now consider another maximal multi-schedule u , and the set of $\langle \text{link}, \text{time slot} \rangle$ pairs that are in $\text{supp}(u)$ but not in the union of the interference neighborhoods of $\langle \text{link}, \text{time slot} \rangle$ pairs in $\text{supp}(s)$, i.e., the set

$$P = \text{supp}(u) \setminus \bigcup_{\langle i, t \rangle \in \text{supp}(s)} I'(i, t).$$

We note that for any $\langle \text{link}, \text{time slot} \rangle$ pair $\langle j, t' \rangle \in P$, we must have $\langle j, \tilde{t} \rangle \in \text{supp}(s)$ for some $\tilde{t} \neq t'$; in other words, link j must be scheduled in s at some time slot other than

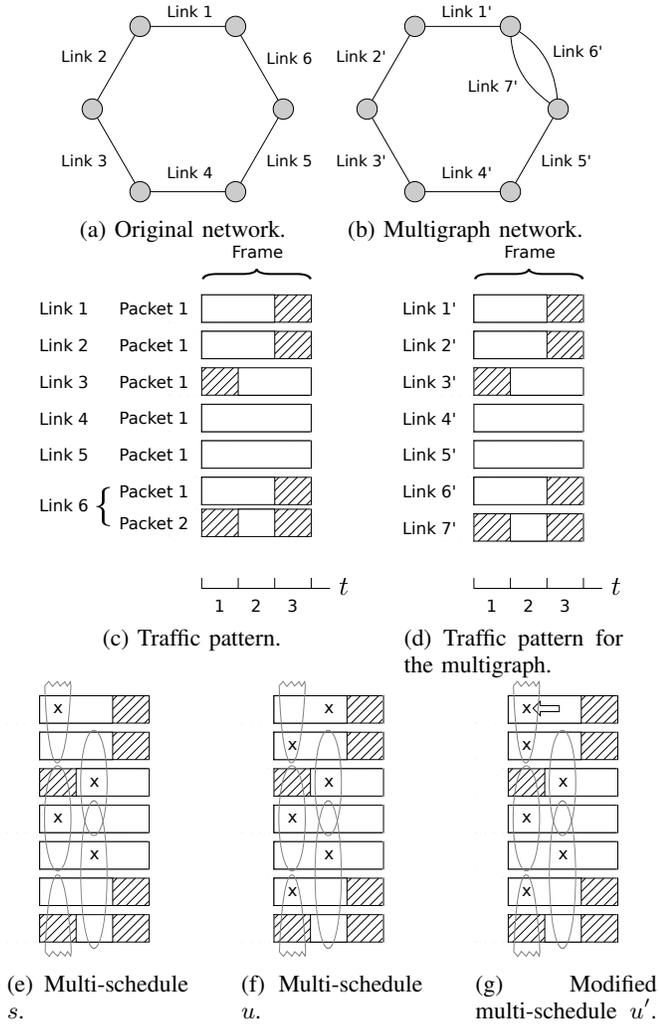


Fig. 12: Consider a six-cycle network with one-hop interference as in (a) and traffic pattern for a frame of 3 time slots as in (c). The network is converted into a multigraph in (b) by dividing the two packets on link 6 into two links with the same end nodes, and the corresponding traffic pattern J is shown in (d). Two maximal multi-schedules, s and u , are given in (e) and (f) with x 's denoting the scheduled links, and the one-hop neighborhoods of the scheduled links of s are illustrated by circles. We see that the one-hop neighborhoods of s cover all scheduled links by s , but miss one link scheduled by u . However, as shown in (g), the missed link can be inserted into the one-hop neighborhood in the first time slot so that all scheduled links by the modified multi-schedule u' are now covered by the one-hop neighborhoods of s .

t' . This holds because otherwise, link j can be added to s at time t without interfering any links in s (note that $\langle j, t' \rangle$ is not in any interference neighborhoods of $\langle \text{link}, \text{time slot} \rangle$ pairs in $\text{supp}(s)$). We use an example in Fig. 12 to illustrate this point.

Let $u|_P$ be the multi-schedule u supported in P ; i.e., $u|_P(i, t) = u(i, t)\mathbb{1}_P(\langle i, t \rangle)$, where

$$\mathbb{1}_P(\langle i, t \rangle) = \begin{cases} 1 & \text{if } \langle i, t \rangle \in P \\ 0 & \text{otherwise} \end{cases}$$

is the indicator function. Then from the analysis above, we know that

$$\|u|_P\| \leq \|s\|.$$

Furthermore, the multi-schedule u can have at most some set of links $I' \in \mathcal{I}'(i)$ active in $I(i, t)$ for any $\langle i, t \rangle \in \text{supp}(s)$ because of the interference. Therefore, we have

$$\begin{aligned} \|u\| &= \|u|_P\| + \|u|_{X \setminus P}\| \\ &\leq \|s\| + \sum_{\langle i, t \rangle \in \text{supp}(s)} \|u|_{I(i, t)}\| \\ &\leq \|s\| + \sum_{\langle i, t \rangle \in \text{supp}(s)} \max_{I \in \mathcal{I}'(i)} \|\mathbb{1}_{I \times \{t\}}\| \\ &= \sum_{\langle i, t \rangle \in \text{supp}(s)} \left(\bar{\alpha}_i + \max_{I \in \mathcal{I}'(i)} \|\mathbb{1}_{I \times \{t\}}\| \right). \end{aligned}$$

Then

$$\begin{aligned} \frac{\|s\|}{\|u\|} &\geq \frac{\sum_{\langle i, t \rangle \in \text{supp}(s)} \bar{\alpha}_i}{\sum_{\langle i, t \rangle \in \text{supp}(s)} (\bar{\alpha}_i + \max_{I \in \mathcal{I}'(i)} \|\mathbb{1}_{I \times \{t\}}\|)} \\ &\geq \min_{\langle i, t \rangle \in \text{supp}(s)} \frac{\bar{\alpha}_i}{\bar{\alpha}_i + \max_{I \in \mathcal{I}'(i)} \|\mathbb{1}_{I \times \{t\}}\|} \\ &\geq \min_{j \in L} \frac{\alpha_j}{\alpha_j + \max_{I \in \mathcal{I}(j, L)} \sum_{k \in I} \alpha_k} \\ &= \psi_L(\alpha) \\ &\geq \psi(\alpha), \end{aligned}$$

where the last inequality comes from Lemma 1. Since s and u are any arbitrary maximal multi-schedules, we have

$$\frac{\min_{\phi \in M_{J, L}} \alpha_L^T \phi}{\max_{\phi \in M_{J, L}} \alpha_L^T \phi} \geq \psi(\alpha).$$

Combining (2) and Lemma 1, we have

$$\begin{aligned} \sigma_L^*(F) &\geq \frac{\min_{\phi \in M_L(F)} \alpha_L^T \phi}{\max_{\phi \in M_L(F)} \alpha_L^T \phi} \\ &= \frac{\sum_{J \in \mathcal{J}(F)} \pi(J) \alpha_L^T \phi_J^{(1)}}{\sum_{J \in \mathcal{J}(F)} \pi(J) \alpha_L^T \phi_J^{(2)}} \\ &\geq \min_{J \in \mathcal{J}(F)} \frac{\alpha_L^T \phi_J^{(1)}}{\alpha_L^T \phi_J^{(2)}} \\ &\geq \psi(\alpha). \end{aligned}$$

By taking supremum over α we get Theorem 2. ■

APPENDIX B PROOF OF THEOREM 3

By the definition of the interference degree β , there exists a link l_0 together with β of its neighbors $\{l_1, l_2, \dots, l_\beta\} \subseteq N(l_0)$ such that no two links in $\{l_1, l_2, \dots, l_\beta\}$ interfere each other. Now let $L = \{l_0, l_1, l_2, \dots, l_\beta\}$, then L induces a star

network with the center link l_0 and β leaves in the conflict graph. In this proof we construct periodic traffic processes with packets only arriving on L , and show that the efficiency ratio of LDF is $\gamma_{\text{LDF}}^* \leq \frac{1}{\sqrt{\beta+1}}$.

Let the traffic be such that packets only arrive on $L = \{l_0, l_1, \dots, l_\beta\}$ at odd time slots with the following 2-patterns (a_i is the number of arriving packets on link l_i , and the vector τ_i are the associated maximum delays)

- Pattern 0 (J_0): $a_i = 1$ for all $i = 0, 1, 2, \dots, \beta$, $\tau_0 = 2$, and $\tau_i = 1$ for $i = 1, 2, \dots, \beta$.
- Pattern j ($J_j, j = 1, 2, \dots, \beta$): $a_0 = 1$, $a_i = 1$, $a_k = 0$ for any other k , $\tau_0 = 1$, $\tau_j = 2$.

Note that since no packets arrive at the even time slots and the maximum delay is 2, all packets expire at the end of the even time slots. Also note that each 2-pattern takes two time slots.

Now consider periodic traffic with a period of $2n$ time slots, where each period consists of n_0 consecutive J_0 's, followed by n_1 consecutive J_1 's, and then n_2 consecutive J_2 's, and so on, till n_β consecutive J_β 's, with $n = \sum_{i=0}^{\beta} n_i$. We then notice that for each traffic pattern J_j , all of the packets can be scheduled if the packets with maximum delay 1 get scheduled in the odd time slots. Hence the optimal scheduler can always achieve the QoS vector $p_{\text{OPT}} = \mathbf{1}$. For LDF, we assume each time the pattern J_j arrives, only deficit j increases by one, and LDF chooses to schedule link j , reducing deficit j by one. Then with all the deficits unchanged, LDF achieves the QoS vector of

$$p_{\text{LDF}} = \left(\frac{n_0}{n}, \frac{n_1}{n_0 + n_1}, \frac{n_2}{n_0 + n_2}, \dots, \frac{n_\beta}{n_0 + n_\beta} \right).$$

Then we can easily see that LDF cannot achieve QoS vector $p = p_{\text{LDF}} + \epsilon \mathbf{1}$ for any $\epsilon > 0$ since adding an extra ϵ/n deficit for each pattern on each link would make the deficits grow without bound. Now let n_0/n approximate $\frac{1}{\sqrt{\beta+1}}$ and let n_j approximate $\frac{1}{\beta + \sqrt{\beta}}$ for $j = 1, 2, \dots, \beta$. Then we will have that p_{LDF} approximates $\frac{1}{\sqrt{\beta+1}} \mathbf{1}$. Therefore the efficiency ratio of LDF under the given traffic and deficit arrival is at most $\frac{1}{\sqrt{\beta+1}}$.

APPENDIX C PROOF OF COROLLARY 2

Choose an arbitrary leaf l_0 in the tree and set it to be the root. Define the depth of each link to be the number of hops it is away from the root. Then the root has depth 0 and any other link has depth greater than 0. For any link l , set the weight of it to be $\alpha_l = (\beta - 1)^{-\text{dep}(l)/2}$, where $\text{dep}(l)$ is the depth of link l . Note that the root has the largest weight 1. Then each link has at most one parent and at most $\beta - 1$ children, where the parent has a weight $\sqrt{\beta - 1}$ times the weight of that link and the children have weights $\frac{1}{\sqrt{\beta - 1}}$ times the weight of that link. Hence the pressure of link l for α is

$$\kappa_l(\alpha) \geq \frac{1}{1 + \sqrt{\beta - 1} + (\beta - 1) \frac{1}{\sqrt{\beta - 1}}} = \frac{1}{2\sqrt{\beta - 1} + 1}.$$

Then by the definition of minimum pressure for α and Theorem 2 we get Corollary 2.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation (NSF) under Grants CNS-1261429, CNS-1264012, and CNS-1262329, in part by the Defense Threat Reduction Agency (DTRA) under Grant HDTRA1-13-1-0030.

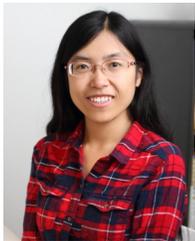
REFERENCES

- [1] I.-H. Hou, V. Borkar, and P. R. Kumar, "A theory of QoS for wireless," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, Rio de Janeiro, Brazil, 2009, pp. 486–494.
- [2] I.-H. Hou and P. R. Kumar, "Admission control and scheduling for QoS guarantees for variable-bit-rate applications on wireless channels," in *Proc. ACM Int. Symp. Mobile Ad Hoc Networking and Computing (MobiHoc)*, New Orleans, LA, 2009, pp. 175–184.
- [3] —, "Scheduling heterogeneous real-time traffic over fading wireless channels," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, San Diego, CA, 2010, pp. 1–9.
- [4] —, "Utility-optimal scheduling in time-varying wireless networks with delay constraints," in *Proc. ACM Int. Symp. Mobile Ad Hoc Networking and Computing (MobiHoc)*, Chicago, IL, 2010, pp. 31–40.
- [5] J. J. Jaramillo and R. Srikant, "Optimal scheduling for fair resource allocation in ad hoc networks with elastic and inelastic traffic," *IEEE/ACM Trans. Netw.*, vol. 19, pp. 1125–1136, Aug. 2011.
- [6] J. J. Jaramillo, R. Srikant, and L. Ying, "Scheduling for optimal rate allocation in ad hoc networks with heterogeneous delay constraints," *IEEE J. Sel. Areas Commun.*, vol. 29, pp. 979–987, May 2011.
- [7] C. Joo, X. Lin, and N. B. Shroff, "Performance limits of greedy maximal matching in multi-hop wireless networks," in *Proc. IEEE Conf. Decision and Control (CDC)*, New Orleans, LA, 2007, pp. 1128–1133.
- [8] A. A. Reddy, S. Sanghavi, and S. Shakkottai, "On the effect of channel fading on greedy scheduling," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, Orlando, FL, 2012, pp. 406–414.
- [9] J. G. Dai, "On positive Harris recurrence of multiclass queueing networks: A unified approach via fluid limit models," *Ann. Appl. Probab.*, vol. 5, no. 1, pp. 49–77, Feb. 1995.
- [10] P. Chaporkar, K. Kar, and S. Sarkar, "Throughput guarantees through maximal scheduling in wireless networks," in *Proc. Annu. Allerton Conf. Communication, Control and Computing*, Monticello, IL, 2005, pp. 28–30.
- [11] A. Dimakis and J. Walrand, "Sufficient conditions for stability of longest queue first scheduling: Second-order properties using fluid limits," *Adv. in Appl. Probab.*, vol. 38, no. 2, pp. 505–521, Jun. 2006.
- [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 Trans. Netw.*, vol. 17, pp. 1132–1145, Aug. 2009.
- [13] R. M. Loynes, "The stability of a queue with non-independent inter-arrival and service times," *Math. Proc. Cambridge Philos. Soc.*, vol. 58, no. 3, pp. 497–520, Jul. 1962.
- [14] S. T. Maguluri, B. Hajek, and R. Srikant, "The stability of longest-queue-first scheduling with variable packet sizes," in *Proc. IEEE Conf. Decision and Control (CDC)*, Orlando, FL, Dec. 2011, pp. 3770–3775.
- [15] R. Srikant and L. Ying, *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. New York: Cambridge Univ. Press, 2014.
- [16] B. Li, C. Boyaci, and Y. Xia, "A refined performance characterization of longest-queue-first policy in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 19, pp. 1382–1395, Oct. 2011.
- [17] —, "Performance guarantee under longest-queue-first schedule in wireless networks," *IEEE Trans. Inf. Theory*, vol. 58, pp. 5878–5889, Sep. 2012.
- [18] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY: Cambridge Univ. Press, 2004.
- [19] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar, and P. Whiting, "Scheduling in a queueing system with asynchronously varying service rates," *Prob. Eng. and Information Sciences*, vol. 18, no. 2, pp. 191–217, Apr. 2004.
- [20] B. Ji, C. Joo, and N. Shroff, "Throughput-optimal scheduling in multihop wireless networks without per-flow information," *IEEE/ACM Trans. Netw.*, vol. 21, pp. 634–647, Apr. 2013.
- [21] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," in *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, San Diego, CA, 2003, pp. 66–80.

- [22] T. Weller and B. Hajek, "Scheduling nonuniform traffic in a packet-switching system with small propagation delay," *IEEE/ACM Trans. Netw.*, vol. 5, pp. 813–823, Dec. 1997.
- [23] J. G. Dai and B. Prabhakar, "The throughput of data switches with and without speedup," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, vol. 2, Tel Aviv, Israel, Mar. 2000, pp. 556–564.
- [24] X. Lin and N. B. Shroff, "The impact of imperfect scheduling on cross-layer rate control in wireless networks," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, vol. 3, Miami, FL, 2005, pp. 1804–1814.
- [25] X. Wu and R. Srikant, "Regulated maximal matching: A distributed scheduling algorithm for multi-hop wireless networks with node-exclusive spectrum sharing," in *Proc. IEEE Conf. Decision and Control (CDC)*, Seville, Spain, 2005, pp. 5342–5347.



Xiaohan Kang received his B.E. degree in electrical engineering from Tsinghua University, Beijing, China, in 2009. He is currently pursuing a Ph.D. degree in electrical engineering at Arizona State University, Tempe, AZ. His research interests include resource allocation in wireless networks and data centers. He received a Schafer 2050 Challenge Graduate Fellowship in 2010 and a University Graduate Fellowship in 2014.



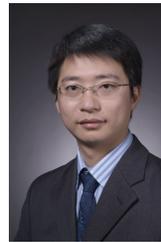
Weina Wang received her B.E. degree in electrical engineering from Tsinghua University, Beijing, China, in 2009. She is currently pursuing a Ph.D. degree in electrical engineering at Arizona State University, Tempe, AZ. Her research interests include resource allocation in communication networks and data privacy.



Juan José Jaramillo received the B.S. degree (summa cum laude) from Universidad Pontificia Bolivariana, Medellín, Colombia, in 1998, and the M.S. and Ph.D. degrees from the University of Illinois at Urbana–Champaign in 2005 and 2009, respectively, all in electrical engineering.

From 2010 to 2011 he was with the Department of Electrical and Computer Engineering at Iowa State University as a Post-Doctoral Research Associate. He is currently with Universidad EAFIT, Colombia, where he is an Assistant Professor with the Department of Applied Math and Engineering. He is the recipient of a Fulbright fellowship.

His research interests include communication networks and game theory.



Lei Ying (M'08) received his B.E. degree from Tsinghua University, Beijing, China, his M.S. and Ph.D. in Electrical and Computer Engineering from the University of Illinois at Urbana–Champaign. He currently is an Associate Professor at the School of Electrical, Computer and Energy Engineering at Arizona State University, and an Associate Editor of the *IEEE/ACM Transactions on Networking*. His research interest is broadly in the area of stochastic networks, including cloud computing, communication networks and social networks. He is coauthor

with R. Srikant of the book *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*, Cambridge University Press, 2014.

He won the Young Investigator Award from the Defense Threat Reduction Agency (DTRA) in 2009 and NSF CAREER Award in 2010. He was the Northrop Grumman Assistant Professor in the Department of Electrical and Computer Engineering at Iowa State University from 2010 to 2012.