Greedy Maximal Matching: Performance Limits for Arbitrary Network Graphs Under the Node-exclusive Interference Model

Changhee Joo, *Member, IEEE*, Xiaojun Lin, *Member, IEEE*, and Ness B. Shroff†, *Fellow, IEEE*

***Abstract*—*Greedy Maximal Matching* (*GMM*) is an important scheduling scheme for multi-hop wireless networks. It is computationally simple, and has often been numerically shown to achieve throughput that is close to optimal. However, to date the performance limits of *GMM* have not been well understood. In particular, although a lower bound on its performance has been well known, this bound has been empirically found to be quite loose. In this paper, we focus on the well-established *node-exclusive* interference model and provide new analytical results that characterize the performance of *GMM* through a topological notion called the *local-pooling factor*. We show that for a given network graph with single-hop traffic, the efficiency ratio of *GMM* (i.e., the worst-case ratio of the throughput of *GMM* to that of the optimal) is equal to its local-pooling factor. Further, we estimate the local-pooling factor for arbitrary network graphs under the node-exclusive interference model and show that the efficiency ratio of *GMM* is no smaller than $\frac{d^*}{2d^*-1}$ in a network topology of maximum node-degree d^*. Using these results, we identify specific network topologies for which the efficiency ratio of *GMM* is strictly less than 1. We also extend the results to the more general scenario with multi-hop traffic, and show that *GMM* can achieve similar efficiency ratios when a flow-regulator is used at each hop.**

I. Introduction

Link scheduling is an important problem in wireless networks [2]–[12]. Since radio transmissions on a common medium can interfere with each other, it is often necessary to schedule transmissions that interfere with each other at different times. While the optimal scheduling algorithm has been well-known [3], such algorithm is of high computational complexity, and is difficult to be deployed in real networks. Recently, several researchers have developed simple scheduling solutions for an important class of interference models called the *node-exclusive* (or *primary*) interference model. Under this interference model, a node cannot simultaneously transmit or receive, and cannot simultaneously communicate with two or more nodes in the network. The node-exclusive model is a good representation of practical wireless systems using Bluetooth or FH-CDMA networks [2], [13], [14]. Under this model, the scheduling problem can be mapped to a *matching* problem, i.e., any active set of links[1] must form a *matching* of the nodes in the network. In this setting, there exists a polynomial-time optimal solution called the *Maximum Weighted Matching* (*MWM*) policy, which finds the matching that maximizes the queue-weighted rate sum. However, the complexity of *MWM* is roughly $O(|V|^3)$ [15], where $|V|$ is the total number of nodes in the network. Hence, it is still too complex to implement in most practical scenarios.

To address this issue, a well-known solution called the *Greedy Maximal Matching* (*GMM*) has been developed that significantly reduces the scheduling complexity [2], [16], [17]. Its schedule is the maximal matching obtained by selecting links in decreasing order of queue-weighted rate. (See Section II for more details.) We can characterize the performance of *GMM* through its efficiency ratio γ^*, which is the largest number γ such that for any offered load $\vec{\lambda}$ that the optimal *MWM* policy can support, *GMM* can support $\gamma\vec{\lambda}$. It is relatively straightforward to show that the efficiency ratio of *GMM* is at least $\frac{1}{2}$ under the node-exclusive interference model, i.e., *GMM* can sustain at least half of the throughput of the optimal *MWM* policy. In fact, simulation results suggest that the performance of *GMM* is often much better than this lower bound in most network settings. Further, it has been shown that if the network topology satisfies the so-called *local pooling* condition [18], [19], then *GMM* can in fact achieve the full capacity region. However, realistic network topologies may not satisfy the local pooling condition, thus sharp results on its efficiency ratio are desirable.

In this paper, our main contribution is to provide new analytical results on the achievable efficiency ratio of *GMM* for a large class of network topologies. Such an evaluation is important for the following reasons: i) It has been empirically observed in [4] that the throughput achieved by *GMM* is close to the maximum achievable throughput in a variety of networking scenarios. ii) *GMM* can be implemented in a distributed manner [20], which is important from the point of view of many multi-hop networking systems. Further, even simpler constant-time-complexity random algorithms can be developed to approximate the performance of *GMM* as shown in [4]. iii) Although many distributed scheduling schemes have

C. Joo and N. B. Shroff are with Departments of ECE and CSE at the Ohio State University, Columbus, Ohio, e-mail: {cjoo, shroff}@ece.osu.edu

X. Lin is with School of Electrical and Computer Engineering at Purdue University, West Lafayette, Indiana, e-mail: linx@ecn.purdue.edu

This work was supported in part by ARO MURI Awards No. W911NF-07-10376 (SA08-03) and W911NF-08-1-0238, and NSF Awards CNS-0626703, CNS-0721236, ANI-0207728, and CCF-0635202, USA, and in part by IT Scholarship Program supervised by IITA and MIC, Republic of Korea.

A preliminary version of this work was presented at IEEE CDC'07 [1].

† Corresponding author. Phone: +1-614-247-6554, Fax: +1-614-292-7596, Email: shroff@ece.osu.edu.

[1] A link is defined as a node pair, for which one of the nodes is transmitter and the other is a receiver. We say that a link is active if a node in this node pair transmits to the other node of the node pair. We assume throughout that the links are bi-directional.

been recently developed [8], [10]–[12], the study of *GMM* continues to remain attractive because, empirically, *GMM* performs better than these schemes [9], either in terms of the achievable throughput, or in terms of the resultant queueing delay.

Along this direction, we generalize the notion of local pooling in [18], and derive an equivalent characterization of the efficiency ratio of *GMM* through a topological property, i.e., the *local-pooling factor*, of the underlying network graph. In particular, we show that the efficiency ratio of *GMM* under a given network topology is equal to its local-pooling factor. We then estimate the local-pooling factor for arbitrary network graphs under the node-exclusive interference model, which enable us to identify network topologies where the efficiency ratio of *GMM* could be much less than 1.

In related work [21], we have studied the Greedy Maximal Scheduling (*GMS*) policy for more general interference models, and have developed an iterative methodology to obtain the bounds on the local-pooling factor. However, the focus of this paper is different, as described in the contributions below:

- We focus on the node-exclusive interference model, under which we are able to derive far sharper bounds on the worst-case efficiency ratio than in [21]. More importantly, we provide these bounds as a function of the maximum node degree, and find network topologies where these bounds are tight.
- This paper allows for arbitrary link capacities, while in [21] all links are assumed to have unit capacity.
- We also consider the scenario with multi-hop traffic. For given multi-hop traffic flows and their paths, we show that *GMM* achieves similar efficiency ratios in multi-hop traffic scenarios as in the single-hop traffic case, when per-flow regulators are used at each hop.
- We explicitly provide the proof of Proposition 1 of [21], which was not presented in [21] because the focus of that paper was on developing a methodology to estimate the local-pooling factor.

The rest of the paper is organized as follows. We first describe our model in Section II. We then introduce the notion of local-pooling factor and show in Section III that the efficiency ratio of *GMM* under an arbitrary network topology is equal to its local-pooling factor. In Section IV, by characterizing the set of unstable links under *GMM*, we estimate the local-pooling factors for arbitrary network graphs under the node-exclusive interference model. These results lead to the discovery of network graphs where the efficiency ratio of *GMM* is low. Finally, we extend the results to the more general multi-hop traffic scenarios in Section V, and we conclude in Section VI.

II. Network Model

We model a wireless network by a graph $G(V,E)$, where V is the set of nodes, and E is the set of undirected links. We assume a time-slotted system, where the length of each time slot is of unit length. Let c_l denote the capacity of each link l, and let the column-vector $\vec{c}^T := [c_1, c_2, \ldots, c_{|E|}]$. We assume that in each time slot, a link can transmit c_l packets provided that the following *node-exclusive interference* constraint is satisfied: if a link l is transmitting data, then no other links that share a common transmitter node or receiver node with link l can transmit at the same time. Hence, any active set of links must form a *matching* of the nodes in V.

Let $\vec{M}_E$ be a *maximal* matching on E, i.e., no more links can be added to $\vec{M}_E$ without violating the node-exclusive interference constraint. We use a vector in $\{0,1\}^{|E|}$ to denote a maximal matching $\vec{M}_E$ such that the k-th element is set to 1 if link $k \in E$ is included in the maximal matching $\vec{M}_E$, and to 0 otherwise. Let $\mathcal{M}_E$ be the set of all possible maximal matchings and let $Co(\mathcal{M}_E)$ denote its convex hull. We also define the set of link rates as $Co(\mathcal{R}_E) := \{\vec{\phi} \mid \vec{\phi} = \vec{c} \circ \vec{\psi} \text{ for } \vec{\psi} \in Co(\mathcal{M}_E)\}$, where the operator denotes component-wise multiplication, i.e., $[x_1, x_2, \ldots] \circ [y_1, y_2, \ldots] := [x_1 y_1, x_2 y_2, \ldots]$.

We first assume single-hop traffic, i.e., packets arrive to each link l and once they are served, they immediately leave the system. This assumption will be relaxed later in Section V. We also assume that the packet arrivals follow a stationary and ergodic process that satisfies the conditions for the fluid limits to exist. One such condition is that the inter-arrival times are *i.i.d.* The other alternate condition is that the number of arrivals at each time-slot are *i.i.d.* across time. Further, we assume that the dynamics of the system can be modeled as a Markov process[2]. Let $A_l(t)$ denote the number of packet arrivals at link l at time slot t. Let $S_l(t)$ denote the number of packets that link l can serve at time slot t. Note that $S_l(t)$ takes the value c_l if link l is activated at time slot t and is zero otherwise. Also let $Q_l(t)$ denote the number of packets queued at link l at the beginning of time slot t. The queue length then evolves according to the following equations:

$$Q_l(t+1) = [Q_l(t) + A_l(t) - S_l(t)]^+, \tag{1}$$

where $[\cdot]^+$ denote the projection to non-negative real numbers.

Definition 1: The network is *stable* if for the queue length process, we have

$$\limsup_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{1}_{\{\sum_{l\in E} Q_l(t) > \eta\}} \to 0, \text{ as } \eta \to \infty, \tag{2}$$

where $\mathbb{1}_{\{\cdot\}}$ denotes an indicator function taking either 0 or 1 based on the occurrence of specified event.

Let λ_l denote the average arrival rate. The *capacity region* (or the *stability region*) under a given scheduling policy is defined as the set of arrival rate vectors $\vec{\lambda} = [\lambda_1, \lambda_2, \ldots, \lambda_{|E|}]$ such that the system is stable. We define the optimal capacity region Λ as the union of the capacity region of all possible stationary scheduling policies. To characterize it, we define a set

$$\begin{aligned} \Lambda' := \{\vec{\lambda} \mid\ & \text{for some } \vec{\phi} \in Co(\mathcal{R}_E), \\ & \lambda_l < \phi_l \text{ if } \phi_l > 0, \text{ and } \lambda_l = 0 \text{ otherwise}\}, \end{aligned} \tag{3}$$

and its closure

$$\bar{\Lambda}' := \{\vec{\lambda} \mid \vec{\lambda} \preceq \vec{\phi}, \text{ for some } \vec{\phi} \in Co(\mathcal{R}_E)\}, \tag{4}$$

where $\vec{x} \preceq \vec{y}$ denotes that $\vec{x}$ is component-wise less than

[2] For example, we can describe the states using queue lengths, remaining interarrival times, and remaining service times under the assumption of *i.i.d.* interarrival and service times [22].

or equal to by $\vec{y}$. The optimal capacity region Λ can be characterized [3] as

$$\Lambda' \subset \Lambda \subset \bar{\Lambda}'. \tag{5}$$

It is well-known that the *Maximal Weighted Matching (MWM)* policy can achieve this optimal capacity region under the node-exclusive interference model. However, its computational complexity ($O(|V|^3)$) is high. In this paper, we are interested in a simpler policy called *Greedy Maximal Matching* (*GMM*). *GMM* operates as follows: At each time slot, it first picks the link l with the largest weight $q_l c_l$, where q_l is the backlog size of link l; it then discards all links that interfere with link l; it then picks the link k with the largest $q_k c_k$ from the remaining links; and this process continues until no links are left. As we discussed in the introduction, in this paper we are interested in characterizing the efficiency ratio of *GMM* under arbitrary network topologies. We formally define the notion of the efficiency ratio as follows.

Definition 2: For a scheduling policy, e.g., *GMM*, we say that it achieves a *fraction* γ *(*$0 \le \gamma \le 1$*) of the capacity region* under a given network topology if it can keep the system stable for any offered load $\vec{\lambda} \in \gamma\Lambda$.

Definition 3: The *efficiency ratio* γ^* of a scheduling policy under a given network topology is the supremum of all γ such that the policy can achieve a fraction γ of the capacity region, i.e.,

$$\begin{aligned}\gamma^* := \sup\{\gamma \mid\ & \text{the system is stable under all offered} \\ & \text{load vectors } \vec{\lambda} \text{ such that } \vec{\lambda} \preceq \gamma\vec{\phi} \\ & \text{for some } \vec{\phi} \in Co(\mathcal{R}_E)\}.\end{aligned} \tag{6}$$

III. An Equivalent Characterization of the Efficiency Ratio of *GMM*

In this section, we derive an equivalent characterization of the efficiency ratio of *GMM* under arbitrary network topologies. We first recall the following definition of *local pooling* from [18]:

Definition 4: Given a network graph $G(V,E)$, a set of links $L \subset E$ satisfies *local pooling*, if there exists a nonzero $\vec{\alpha} \in \mathbb{R}_+^{|L|}$ such that $\vec{\alpha}^T\vec{\phi}$ is a positive constant for all $\vec{\phi} \in Co(\mathcal{R}_L)$. The graph $G(V,E)$ satisfies *local pooling* if every $L \subset E$ satisfies local pooling.

An example of graphs that satisfy local pooling is the triangular network topology with three nodes and three links of unit capacity as shown in Fig. 1. In this graph, we have three maximal matchings; $[1,0,0]$, $[0,1,0]$, and $[0,0,1]$. For any convex combination $\vec{\phi}$ of these three vectors, we have $\vec{\alpha}^T\vec{\phi} = 1$ with $\vec{\alpha} = [1,1,1]$.

Note that if a set of links L satisfies local pooling, no vector in $Co(\mathcal{R}_L)$ is strictly (component-wise) greater than another vector in $Co(\mathcal{R}_L)$[3]. Dimakis and Walrand [18] have shown that if a network graph satisfies local pooling, *GMM* achieves the full capacity region.

[3]We can prove this by contradiction. Suppose that there exist $\vec{\phi}_1, \vec{\phi}_2 \in Co(\mathcal{R}_L)$ such that $\vec{\phi}_1 \succ \vec{\phi}_2$. Multiplying $\vec{\alpha}$ to both sides, we obtain $\vec{\alpha}^T\vec{\phi}_1 > \vec{\alpha}^T\vec{\phi}_2$, which contradicts the assumption.

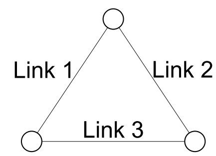

Fig. 1. Triangular network topology with three nodes and three links.

In this paper, we are interested in arbitrary network topologies that may not satisfy local pooling. We now generalize the notion of local pooling to that of the local-pooling factor.

Definition 5: A set of links L satisfies σ*-local pooling*, if $\sigma\vec{\mu} \not\succeq \vec{\nu}$ for all $\vec{\mu}, \vec{\nu} \in Co(\mathcal{R}_L)$. In other words, for all $\vec{\mu}, \vec{\nu} \in Co(\mathcal{R}_L)$, there must exist some $k \in L$ such that $\sigma\mu_k < \nu_k$.

Note that σ cannot be greater than 1 since we can take $\vec{\mu} = \vec{\nu}$. In addition, if a graph, e.g., the triangular network topology, satisfies local pooling of Definition 4, then it must satisfy σ-local pooling for any $\sigma < 1$. We can prove this by contradiction. Suppose that there exist two convex combinations $\vec{\phi}_1, \vec{\phi}_2$ and $\sigma < 1$ such that $\sigma\vec{\phi}_1 - \vec{\phi}_2 \succeq \vec{0}$. Since the graph satisfies local pooling, there exists an $\vec{\alpha}$ such that $\vec{\alpha}^T\vec{\phi}_1 = \vec{\alpha}^T\vec{\phi}_2 > 0$. Multiplying $\vec{\alpha}$ on both sides of $\sigma\vec{\phi}_1 - \vec{\phi}_2 \succeq \vec{0}$, we obtain $\sigma - 1 \ge 0$, which contradicts the assumption.

Definition 6: The *local-pooling factor* of a graph $G(V,E)$ is the supremum of all σ such that every subset $L \in E$ satisfies σ-local pooling. In other words,

$$\begin{aligned}\sigma^* :=& \sup\{\sigma \mid \sigma\vec{\mu} \not\succeq \vec{\nu} \text{ for all } L \text{ and all } \vec{\mu},\vec{\nu} \in Co(\mathcal{R}_L)\} \\ =& \inf\{\sigma \mid \sigma\vec{\mu} \succeq \vec{\nu} \text{ for some } L \text{ and some } \vec{\mu},\vec{\nu} \in Co(\mathcal{R}_L)\}.\end{aligned} \tag{7}$$

By definition, if the local-pooling factor of a graph is σ^*, then every subset $L \subset E$ must satisfy σ^*-local pooling. Note that Definition 4 of local pooling corresponds to $\sigma^* = 1$. The results of [18] imply that if the local-pooling factor of the graph is 1, then the efficiency ratio of *GMM* will be 1. We next generalize this result to the case when $\sigma^* < 1$. We start with two lemmas.

Lemma 1: If the local-pooling factor of a graph $G(V,E)$ is σ^*, then the efficiency ratio γ^* of *GMM* under this network topology is no smaller than σ^*, i.e., $\gamma^* \ge \sigma^*$.

Proof: We need to show that for any offered load $\vec{\lambda}$ strictly within $\sigma^*\Lambda$, the network is stable under *GMM*. We prove stability by finding a Lyapunov function with negative drift for the fluid limit model of the system.

We first define the fluid limit model of the system as in [5], [22]. We recall that $A_l(t)$ and $S_l(t)$ denote the number of packet arrivals and available service, respectively, at time slot t, and that $Q_l(t)$ denotes the number of queued packets at the beginning of time slot t and it evolves according to (1). We interpolate the values of $A_l(t)$ and $S_l(t)$ to all nonnegative real numbers t by setting $A_l(t) = A_l(\lfloor t \rfloor)$ and $S_l(t) = S_l(\lfloor t \rfloor)$, where $\lfloor t \rfloor$ denotes the largest integer smaller than or equal to t. We also interpolate the values of $Q_l(t)$ by linear interpolation between $\lfloor t \rfloor$ and $\lfloor t \rfloor + 1$. Then, using the techniques of Theorem 4.1 of [22], we can show that, for almost all sample paths and for all positive sequence $x_n \to \infty$, there exists a subsequence x_{n_j} with $x_{n_j} \to \infty$ such that the following convergence holds uniformly over compact intervals of time t: For all $l \in E$, there exist limits λ_l, $q_l(t)$, and $\pi_l(t)$ such that

$$\begin{aligned} \tfrac{1}{x_{n_j}} \int_0^{x_{n_j} t} A_l(s) ds &\to \lambda_l t, \\ \tfrac{1}{x_{n_j}} \int_0^{x_{n_j} t} S_l(s) ds &\to \int_0^t \pi_l(s) ds, \\ \tfrac{1}{x_{n_j}} Q_l(x_{n_j} t) &\to q_l(t). \end{aligned} \tag{8}$$

Moreover, for all $l \in E$, the limits $q_l(t)$ and $\pi_l(t)$ satisfy

$$\frac{d}{dt} q_l(t) = \begin{cases} \lambda_l - \pi_l(t), & \text{if } \lambda_l - \pi_l(t) \geq 0, \text{ or } q_l(t) > 0, \\ 0, & \text{otherwise,} \end{cases} \tag{9}$$

and $\vec{\pi}(t)$ should be a convex combination of maximal matchings chosen by *GMM*, which will be further explained in the following. Any such limit $[\vec{q}(t), \vec{\pi}(t)]$ is called a *fluid limit* of the system.

We now use the idea from [18] and show that for any offered load λ strictly within $\sigma^* \Lambda$, the largest queue-weighted rate $\max_l q_l c_l$ of the fluid limit model always decreases under *GMM*. Note that $q_l(t)$ is absolutely continuous, and hence its derivative exists almost everywhere. Consider those times t when the derivative $\frac{d}{dt} q_l(t)$ exists for all $l \in E$. Let $L_0(t)$ denote the set of links with the largest queue-weighted rate at time t, i.e.,

$$L_0(t) := \left\{ l \in E \;\middle|\; q_l(t) c_l = \max_{k \in E} q_k(t) c_k \right\}.$$

Let $L(t)$ denote the set of links with the largest derivative of the queue-weighted rate among the links in $L_0(t)$,

$$L(t) := \left\{ l \in L_0(t) \;\middle|\; \frac{d}{dt} q_l(t) c_l = \max_{k \in L_0(t)} \frac{d}{dt} q_k(t) c_k \right\}.$$

Next, we show by contradiction that there cannot exist a link $l \in L(t)$ such that $\frac{d}{dt} q_l(t) c_l \geq 0$ and $q_l(t) > 0$.

Suppose that we have $\vec{\pi}(t)$ as a result of scheduling under *GMM*, and there exists a link $l \in L(t)$ with $\frac{d}{dt} q_l(t) c_l \geq 0$ and $q_l(t) > 0$. Then by the definition of $L(t)$, all links in $L(t)$ have the same non-negative derivative. Since links in $L(t)$ have the largest derivative among links in $L_0(t)$ at time t and $q_l(t)$'s are continuous, there exists a small $\delta > 0$ such that for all $\tau \in (t, t + \delta]$,

$$\min_{l \in L(t)} q_l(\tau) c_l > \max_{k \in E \setminus L(t)} q_k(\tau) c_k.$$

Since *GMM* finds a maximal matching in decreasing order of queue-weighted rate, links in $L(t)$ should be chosen first during the interval $(t, t + \delta]$. Hence, each schedule of *GMM*, when projected to the set $L(t)$, should be a maximal matching on $L(t)$, and thus, the average service rate $\vec{\pi}^*$ during $(t, t+\delta]$ should satisfy that $\vec{\pi}^*|_{L(t)} \in Co(\mathcal{R}_{L(t)})$, where $\cdot|_{L(t)}$ denote the projection of a vector onto $L(t)$. Then, by letting $\delta \to 0$, we obtain $\vec{\pi}^* \to \vec{\pi}(t)$, and hence, the limit must satisfy that

$$\vec{\pi}(t)|_{L(t)} \in Co(\mathcal{R}_{L(t)}). \tag{10}$$

We refer readers to [23] for the complete proof of this part.

Since the local-pooling factor is σ^*, and $\vec{\lambda}$ falls strictly within $\sigma^* \Lambda$ (i.e., $\vec{\lambda} \in \sigma^* \Lambda'$), there must exist a $k \in L(t)$ such that $\lambda_k < \pi_k(t)$. Define

$$\epsilon_* := \inf_{\vec{\phi} \in Co(\mathcal{R}_L), L \subset E} \left\{ \max_{k \in L} (\phi_k - \lambda_k) \right\}. \tag{11}$$

Note that since $\vec{\lambda}$ falls strictly within $\sigma^* \Lambda$ and the local pooling factor is σ^*, we must have $\epsilon_* > 0$. Hence, by the earlier argument, there exists a $k \in L(t)$ such that $\lambda_k - \pi_k(t) \leq -\epsilon_*$. This implies that $\frac{d}{dt} q_l(t) c_l = \frac{d}{dt} q_k(t) c_k \leq -\epsilon_*$ for all $l \in L(t)$. This result contradicts our assumption and implies that the largest queue-weighted rate must decrease.

Therefore, we can pick the Lyapunov function as $V(t) := \max_{l \in E} q_l(t) c_l$. We have, if $V(t) > 0$,

$$\frac{D^+}{dt^+} V(t) \leq \max_{l \in L_0(t)} \frac{d}{dt} q_l(t) c_l = \left. \frac{d}{dt} q_k(t) c_k \right|_{k \in L(t)} \leq -\epsilon_*,$$

where $\frac{D^+}{dt^+} V(t) = \lim_{\delta \downarrow 0} \frac{V(t+\delta) - V(t)}{\delta}$. Since the above inequality holds for almost every t, it implies that the fluid limit model of the system is stable. Hence, by Theorem 4.2 of [22], the original system is also stable. ■

Lemma 1 shows that the efficiency ratio of *GMM* under an arbitrary network graph is no smaller than the local-pooling factor, i.e., $\gamma^* \geq \sigma^*$.

The next lemma shows that $\gamma^* \leq \sigma^*$.

Lemma 2: If there exist a subset of links $L \subset E$, a positive number σ, and two vectors $\vec{\mu}, \vec{\nu} \in Co(\mathcal{R}_L)$ such that $\sigma \vec{\mu} \succeq \vec{\nu}$, then, for arbitrarily small $\epsilon > 0$, there exists a traffic pattern with offered load $\vec{\nu} + \epsilon \vec{e}_L$ such that the system is unstable under *GMM*, where $\vec{e}_L$ is a vector with $e_l = \frac{1}{c_l}$ for $l \in L$ and $e_l = 0$ for $l \notin L$.

Remark: Since $\vec{\nu} \in \sigma \Lambda$, Lemma 2 implies that the efficiency ratio of *GMM* under this network topology is no greater than the local-pooling factor, i.e., $\gamma^* \leq \sigma^*$.

Proof: We will construct a traffic pattern with offered load $\vec{\nu} + \epsilon \vec{e}_L$ based on $\vec{\nu}$, and show that under this traffic pattern, the queue length will increase to infinity under *GMM*.

Let $\aleph$ denote the number of all maximal matchings $\vec{M}_i$ on L. Since $\vec{\nu}$ is a convex combination of these maximal matchings, it can be written as

$$\vec{\nu} = \sum_{i=0}^{\aleph - 1} w_i (\vec{c} \circ \vec{M}_i), \tag{12}$$

where $w_i \geq 0$ for all $0 \leq i \leq \aleph - 1$ and $\sum_{i=0}^{\aleph-1} w_i = 1$.

We now construct a traffic pattern with offered load $\vec{\lambda} = \vec{\nu} + \epsilon \vec{e}_L$ such that the system is unstable under *GMM*. We assume that packets arrive to a link *before* a time slot and that the queue of all links in L is empty at the beginning.

Let $\tilde{c}$ denote the least common multiple of all c_l's. At each time slot t, pick a matching $\vec{M}_i$ with probability w_i for $i = 0, 1, \ldots, \aleph - 1$. Then *with probability* $(1 - \epsilon')$, we inject c_l packets to links included in $\vec{M}_i$, and *with probability* ϵ', inject $c_l + \frac{\tilde{c}}{c_l}$ packets to links l included in $\vec{M}_i$ and $\frac{\tilde{c}}{c_k}$ packets to all other links k in L. The arrival pattern repeats at every time slot.

We can show by induction that all links in L have the same queue-weighted rates that keep increasing. Suppose that all links in L have the same queue-weighted rates at the end of time slot $t-1$. (Note that this is true at $t=1$ if the system is initially empty.) At time slot t, if we inject c_l packets to links l in some maximal matching $\vec{M}_j$, then since the links with new packet arrivals have the largest queue-weighted rates and they do not interfere with each other, these links will be served simultaneously under *GMM* during the time slot. Hence, at the end of time slot t, all queue-weighted rates in L will remain the same. On the other hand, if we inject $c_l + \frac{\tilde{c}}{c_l}$ packet to links l included in some maximal matching $\vec{M}_j$ and $\frac{\tilde{c}}{c_k}$ to all other links k in L, the links in $\vec{M}_j$ again have the largest queue-weighted rates and will be served simultaneously under *GMM*. As a result, at the end of the time slot, all links in L have the same queue-weighted rates, which however increase by $\tilde{c}$ from those in the previous time slot. Since $\vec{M}_j$ is an arbitrary maximal matching on L, we can conclude that, at the end of each time slot, all queue-weighted rates in L will be the same. However, with probability ϵ', the queue length of each link $l \in L$ increases by $\frac{\tilde{c}}{c_l}$.

The average arrival rate of this traffic pattern can be estimated as

$$\sum_{i=0}^{\aleph-1} w_i(1-\epsilon')\cdot\vec{c}\circ\vec{M}_i + \sum_{i=0}^{\aleph-1} w_i\epsilon'\cdot(\vec{c}\circ\vec{M}_i + \tilde{c}\vec{e}_L) = \vec{\nu} + \epsilon\vec{e}_L,$$

where $\epsilon = \epsilon'\tilde{c}$. Hence, we have shown that the system with offered load $\vec{\nu} + \epsilon\vec{e}_L$ is unstable under *GMM*. ■

Note that the key to the proof is to construct a traffic pattern such that (i) it keeps all links in L of the same queue-weighted rate, and (ii) it injects packets according to the maximal matchings that form the vector $\vec{\nu}$ so that these maximal matchings will be picked by *GMM*. In fact, for any vector $\vec{\nu} \in Co(\mathcal{R}_L)$, we can construct a traffic pattern with $\vec{\nu} + \epsilon\vec{e}_L$ such that the system is unstable under *GMM*. The following example shows that such a vector $\vec{\nu}$ is not necessarily on the boundary of the optimal capacity region.

Example: The following example illustrates how such a traffic pattern can be constructed in the 6-link cycle network shown in Fig. 2. Assume that all links have unit capacity. We number all links clockwise from 0 to 5. All possible maximal matchings under this network graph are listed below.

- $\vec{M}_0 = [1,0,1,0,1,0]$, $\vec{M}_1 = [0,1,0,1,0,1]$,
- $\vec{M}_2 = [1,0,0,1,0,0]$, $\vec{M}_3 = [0,0,1,0,0,1]$, $\vec{M}_4 = [0,1,0,0,1,0]$.

Note that the number of links included in a maximal matching is three for $\vec{M}_0$ and $\vec{M}_1$, and is two for $\vec{M}_2$, $\vec{M}_3$, and $\vec{M}_4$. Figs. 2(b) and 2(c) show the two instances of the maximal matchings, i.e., $\vec{M}_0$ and $\vec{M}_2$. Note that if we choose two vectors $\vec{\mu}, \vec{\nu}$ from the convex set of maximal matchings $Co(\{\vec{M}_i\})$ as

$$\vec{\mu} = \tfrac{1}{2}\vec{M}_0 + \tfrac{1}{2}\vec{M}_1 = \left[\tfrac{1}{2},\tfrac{1}{2},\tfrac{1}{2},\tfrac{1}{2},\tfrac{1}{2},\tfrac{1}{2}\right],$$
$$\vec{\nu} = \tfrac{1}{3}\vec{M}_2 + \tfrac{1}{3}\vec{M}_3 + \tfrac{1}{3}\vec{M}_4 = \left[\tfrac{1}{3},\tfrac{1}{3},\tfrac{1}{3},\tfrac{1}{3},\tfrac{1}{3},\tfrac{1}{3}\right],$$

then $\frac{2}{3}\vec{\mu} \succeq \vec{\nu}$.

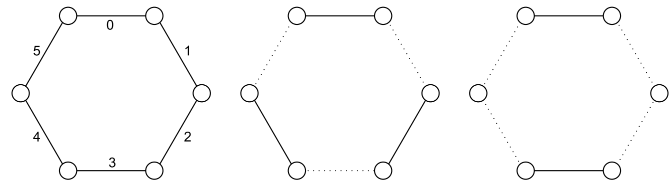

(a) Topology (b) Maximal matching $\vec{M}_0$ (c) Maximal matching $\vec{M}_2$

Fig. 2. The 6-link cycle network and the instances of maximal matching. The solid lines in (b) and (c) are the active links.

We now construct a traffic pattern with offered load $\vec{\lambda} = \vec{\nu} + \epsilon\vec{e}$ such that the system is unstable under *GMM*, where $\vec{e} = [1,1,1,1,1,1]$ and ϵ is a small positive number. Assume that all queues in the system are of the same length at time 0.

At each time slot, packets are injected to links in L as follows: With probability $\frac{1}{3}$ each, pick a maximal matching from $\vec{M}_2$, $\vec{M}_3$, and $\vec{M}_4$. Then,

- *with probability* $(1-\epsilon)$, one packet is injected to links in the maximal matching.
- *with probability* ϵ, two packets are injected to links in the maximal matching and one packet to all other links, resulting in the overall packet arrivals of $\vec{M}_2+\vec{e}$, $\vec{M}_3+\vec{e}$, and $\vec{M}_4+\vec{e}$, respectively.

Over all links, the arrival rate is $\frac{1}{3}+\epsilon$ and the queue length increases by 1 with probability ϵ every time slot. Hence, the system with offered load $\vec{\nu}+\epsilon$ is unstable under *GMM*. However, the optimal *MWM* policy can support the offered load $\vec{\mu} = \frac{3}{2}\vec{\nu}$ in this example. Hence, the efficiency ratio of *GMM* is no greater than $\frac{2}{3}$ in this 6-link cycle network.

From Lemmas 1 and 2, it directly follows that:

Proposition 3: The efficiency ratio γ^* of *GMM* under a given network topology is equal to its local-pooling factor σ^*.

This result provides an equivalent characterization of the efficiency ratio of *GMM* through the topological properties (i.e., the local-pooling factor) of the given graph. Unfortunately, it can still be quite difficult to compute the local-pooling factor for an arbitrary network graph. We next estimate the local-pooling factors for arbitrary network graphs under the node-exclusive interference model.

IV. Estimates of the Local-Pooling Factor for Arbitrary Network Graphs

In this section, we would like to answer the following questions: (i) how do we estimate the local-pooling factor of a given graph? and (ii) what types of graphs will have low local-pooling factors? We now argue that both questions are intimately related to the characterization of the possible sets of unstable links. Note that in order to claim $\sigma^* \leq \sigma$, we must find a subset of links L, and two vectors $\vec{\mu}, \vec{\nu} \in Co(\mathcal{R}_L)$ such that $\sigma\vec{\mu} \succeq \vec{\nu}$. In fact, in the proof of Lemma 2, we show that for any $\epsilon > 0$, there exists a traffic pattern with offered load $\vec{\nu} + \epsilon\vec{e}_L$ such that the queues of all links in L increase to infinity together under *GMM*. Hence, a starting point to search for $\vec{\mu}$ and $\vec{\nu}$ would be to find the subset of links L that have queue lengths increasing to infinity together under *GMM* at such an offered load $\vec{\lambda} = \vec{\nu} + \epsilon\vec{e}_L$ and under such a traffic pattern.

To avoid confusion, we let Y denote the set of links in E whose queue lengths increase to infinity together under *GMM* at offered load $\vec{\lambda} = \vec{\nu} + \epsilon\vec{e}_Y$, where $\vec{\nu} \in Co(\mathcal{R}_Y)$. By constructing the traffic pattern as in (1)-(3) earlier in the proof of Lemma 2, we have $Q_l(t)c_l = 0$ for all $l \notin Y$, and there exists a sequence $\tilde{Q}(t)$ such that $Q_l(t)c_l = \tilde{Q}(t)$ for all $l \in Y$ and $\tilde{Q}(t_k) \to \infty$ for a subsequence $\{t_1, t_2, \ldots, \infty\}$. We refer to the links in Y as the *unstable links*. Let X denote the set of nodes connected to any of the links in Y. We call the graph $U(X,Y)$ an *unstable subgraph* of $G(V,E)$, We next define the notion of an *isolated unstable* link and an *open unstable* link in the unstable subgraph $U(X,Y)$.

Definition 7: A link $l \in Y$ connecting two nodes n_1 and n_2 is an *isolated unstable* link if both n_1 and n_2 are of degree 1 in the unstable subgraph $U(X,Y)$.

Definition 8: A link $l \in Y$ connecting two nodes n_1 and n_2 is an *open unstable* link if either n_1 or n_2 is of degree 1 in the unstable subgraph $U(X,Y)$.

We have the following two results.

Lemma 4: If $\vec{\lambda} = \vec{\nu} + \epsilon\vec{e}_Y$ is strictly within Λ, then under *GMM* there is no isolated unstable link in Y.

Proof: Suppose that Y includes an isolated link l. By assumption, link l has no neighboring links in Y and should be included in all maximal matchings on Y. As a result, link l will be selected at all time slots by *GMM*. Since $\lambda_l < c_l$, the queue length of link l cannot increase to infinity. This contradicts the assumption that link l is unstable. ■

Lemma 5: If $\vec{\lambda} = \vec{\nu} + \epsilon\vec{e}_Y$ is strictly within Λ, there is no open unstable link in Y under *GMM*.

Proof: Suppose that Y includes an open unstable link $l_0 = (n_1, n_2)$. Without loss of generality, assume that node n_1 is shared by other unstable links $\{l_1, l_2, \ldots, l_i\} \subset Y$, and node n_2 is of degree 1 in Y.

Note that every maximal matching on Y should include at least one of the links $l_0, l_1, \ldots, l_i$; because, if none of the links $l_1, \ldots, l_i$ are included, link l_0 should then be included in order for the matching to be maximal in Y. Hence, under *GMM*, the sum of the fraction of time that all of these links are served is $\sum_{k\in\{l_0,l_1,\ldots,l_i\}} \frac{\pi_k}{c_k} = 1$. Recall that all queue-weighted rate $Q_{l_i}c_{l_i}$ of links $l_0, l_1, \ldots, l_i$ are the same at $t = t_1, t_2, \ldots$. Since $\sum_{k\in\{l_0,l_1,\ldots,l_i\}} \frac{\lambda_k}{c_k} < 1$, these links cannot be unstable. This contradicts the assumption that the queues of these links increase to infinity together. ■

The above two lemmas imply that any link in Y must belong to a cycle formed by links in Y. Note that it immediately gives us the result that *GMM* achieves the full capacity region in tree networks [18], [19].

In the following lemma, we characterize the property of the unstable subgraph when the arrival rate $\vec{\lambda}$ is within $\gamma\Lambda$.

Lemma 6: Suppose that $\gamma \in (1/2, 1]$ and that $\vec{\lambda} = \vec{\nu} + \epsilon\vec{e}_Y$ is strictly within $\gamma\Lambda$, then the degree of every node $v \in X$ in the unstable subgraph $U(X,Y)$ must be larger than $\frac{\gamma}{2\gamma-1}$.

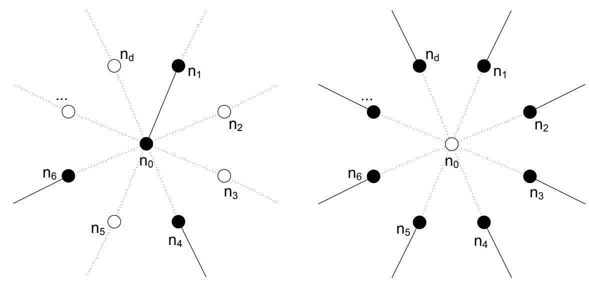

(a) A case that node n_0 is black

(b) Case that node n_0 is white

Fig. 3. Maximal matchings on an unstable network with a degree-d node n_0. Links are denoted by solid lines (when active) and dotted lines (when inactive).

Proof: We consider a node $n_0 \in X$ of degree d (in X) with neighbors $\{n_1, n_2, \ldots, n_d\} \subset X$. Let l_i denote link (n_0, n_i) and let $\mathcal{L}_i$ denote the set of unstable links connected to n_i, excluding l_i, i.e., $\mathcal{L}_i = Y \cap \mathcal{E}(n_i)\backslash\{l_i\}$, where $\mathcal{E}(n_i) \subset E$ is the set of links that are connected to node n_i. In the sequel, we study the activity of node n_0, which can be interpreted as the sum of link activities of those connected to n_0. We focus on the node activity because the optimal capacity region can be characterized by a set of constraints on the node activities under the node-exclusive interference model (see the set Ψ defined below). Through these constraints, we can understand the property of the unstable subgraph.

Observe that all maximal matchings on Y must fall into one of the following two cases:

1) A maximal matching on Y includes a link l_i. In this case, we say that node n_0 is black (see Fig. 3(a)).
2) A maximal matching on Y includes a link from each $\mathcal{L}_i$. In this case, we say that node n_0 is white (see Fig. 3(b)).

We first show that the fraction of time that n_0 is black (the first case) is no greater than γ. Let $A_n := \sum_{l\in\mathcal{E}(n)} \frac{\lambda_l}{c_l}$ denote the weighted arrival rate at node n, and let $S_n := \sum_{l\in\mathcal{E}(n)} \frac{\pi_l}{c_l}$ denote the time-average of the weighted service rate at node n, where π_l is the time-average of the service rate at link l. Note that the optimal capacity region Λ is bounded by

$$\Lambda \subset \Psi := \left\{ \vec{\lambda} \;\middle|\; \sum_{l\in\mathcal{E}(n)} \frac{\lambda_l}{c_l} \leq 1, \text{ for all } n \in E \right\}.$$

By assumption $\vec{\lambda} \in \gamma\Lambda$, we have

$$\sum_{l\in\mathcal{E}(n_0)\cap Y} \frac{\lambda_l}{c_l} \leq \sum_{l\in\mathcal{E}(n_0)} \frac{\lambda_l}{c_l} = A_{n_0} \leq \gamma. \tag{13}$$

If the fraction of time that n_0 is black is greater than γ, then the arrival rate at node n_0 will be smaller than the service rate at n_0, which implies that the queues at the links incident to node n_0 cannot increase to infinity together. This contradicts our assumption.

We next count the total (weighted) service rates over all nodes $n_0, n_1, \ldots, n_d$. Let β denote the fraction of time that node n_0 is black, $0 \leq \beta \leq \gamma$. If node n_0 is black, then at least two nodes (one is n_0) are served. If node n_0 is white, then all nodes $\{n_1, n_2, \ldots, n_d\}$ are served. Hence, we have

$$\sum_{k=0}^{d} S_{n_k} \geq 2\beta + d(1-\beta) \geq 2\gamma + d(1-\gamma). \quad (14)$$

In the last inequality, we have used $0 \leq \beta \leq \gamma$ and $d \geq 2$ (by Lemma 5).

Using the assumption that $\vec{\lambda}$ falls strictly in $\gamma\Psi$, we have

$$\sum_{k=0}^{d} A_{n_k} < \gamma(d+1). \quad (15)$$

We must have

$$\sum_{k=0}^{d} S_{n_k} \leq \sum_{k=0}^{d} A_{n_k}, \quad (16)$$

since, otherwise, the queue lengths of these links cannot increase to infinity together. Combining (14), (15), and (16), we obtain

$$d > \frac{\gamma}{2\gamma - 1}. \quad (17)$$

■

The above lemma immediately implies the second main result of the paper.

Proposition 7: For a given network graph $G(V, E)$ where the largest node degree is d^*, the efficiency ratio γ^* of *GMM* must be no smaller than $\frac{d^*}{2d^*-1}$.

Proof: Suppose that the efficiency ratio is smaller than $\frac{d^*}{2d^*-1}$. Then, according to Proposition 3, we have $\sigma^* < \frac{d^*}{2d^*-1}$. Hence, from Definition 6, there must exist a subset $L \subset E$ and $\vec{\mu}, \vec{\nu} \in Co(\mathcal{R}_L)$ such that $\sigma\vec{\mu} \geq \vec{\nu}$ for some $\sigma < \frac{d^*}{2d^*-1}$. Using Lemma 2, there exists a traffic pattern with $\vec{\lambda} = \vec{\nu} + \epsilon\vec{e}_L$, such that the queue lengths of links in L increase to infinity together. By choosing ϵ small, we can have $\vec{\lambda}$ fall strictly in $\frac{d^*}{2d^*-1}\Lambda$. Then, using Lemma 6, the degree of every node in the unstable graph must be larger than d^*. This contradicts the assumption that the largest node-degree is d^*. ■

Note that the results of Proposition 7 cannot be directly extended to more general K-hop interference models, under which two links within K-hop distance cannot transmit simultaneously. This is because the development of Lemma 6 is based on the node-exclusive interference constraints, and it appears to be difficult to generalize to the K-hop interference models. We refer the readers to [21] for techniques developed for estimating the efficiency ratio of the greedy maximal scheduling policy under the K-hop interference models.

According to Proposition 7, in order to find network topologies where the efficiency ratio of *GMM* is low, we must look at those graphs where the maximum node-degree is high. We have been able to find such graphs where the bound in Proposition 7 is tight with $d^* = 2$ and $d^* = 3$.

A. An example network scenario with $d^* = 2$ and $\gamma^* = \frac{2}{3}$

We consider graphs with degree two. If the graph is a line, then by Lemma 5, *GMM* achieves the full capacity region. Let us instead consider the case when the graph forms a cycle. In the proof of Lemma 2, we show an example of a 6-link cycle network, which has $\gamma^* \leq \frac{2}{3}$. Since this graph has a maximum node-degree of two, Lemma 6 implies that $\gamma^* \geq \frac{2}{3}$. Therefore, *GMM* has an efficiency ratio $\gamma^* = \frac{2}{3}$ in the 6-link cycle network. To the best of our knowledge, *this is the first result that provides the* exact *efficiency ratio for a network graph where GMM cannot achieve the full capacity region.*

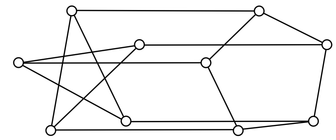

Fig. 4. Star-pentagon Topology

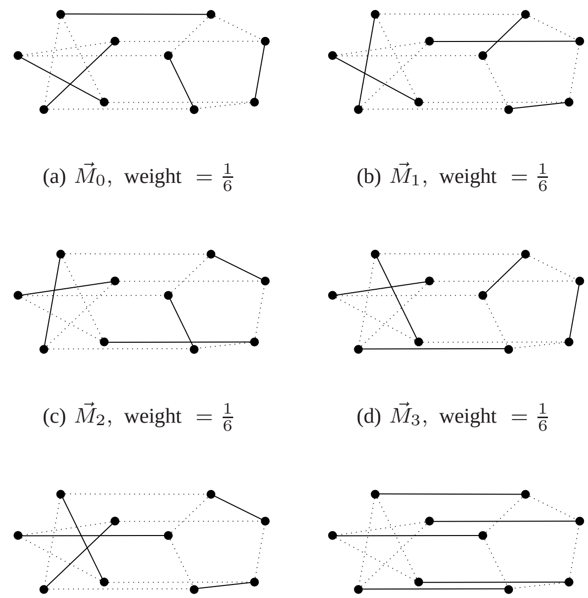

(a) $\vec{M}_0$, weight $= \frac{1}{6}$ (b) $\vec{M}_1$, weight $= \frac{1}{6}$

(c) $\vec{M}_2$, weight $= \frac{1}{6}$ (d) $\vec{M}_3$, weight $= \frac{1}{6}$

(e) $\vec{M}_4$, weight $= \frac{1}{6}$ (f) $\vec{M}_5$, weight $= \frac{1}{6}$

Fig. 5. Maximal matchings for constructing $\vec{\mu}$.

B. An example network scenario with $d^* = 3$ and $\gamma^* = \frac{3}{5}$

We consider the graph with node-degree three as shown in Fig. 4, where all links have unit capacity. We now find two vectors $\vec{\mu}, \vec{\nu} \in Co(\mathcal{M}_E)$ such that $\frac{3}{5}\vec{\mu} = \vec{\nu}$. Fig. 5 shows six maximal matchings and their corresponding weights. The solid lines indicate active links. We choose vector $\vec{\mu}$ as a combination of these matchings, i.e., $\vec{\mu} = \sum_{i=0}^{5}\left(\frac{1}{6}\vec{M}_i\right)$. Fig. 6 illustrates another set of maximal matchings. We choose $\vec{\nu}$ using these matchings as $\vec{\nu} = \sum_{j=6}^{10}\left(\frac{1}{5}\vec{M}_j\right)$. Note that $\mu_l = \frac{1}{3}$ and $\nu_l = \frac{1}{5}$ for all links l.

Since $\frac{3}{5}\vec{\mu} = \vec{\nu}$, the local-pooling factor σ^* cannot be greater than $\frac{3}{5}$, which implies that the efficiency ratio of *GMM* is no greater than $\frac{3}{5}$. However, since the node degree is 3, Proposition 7 implies that the efficiency ratio is no smaller than $\frac{3}{5}$. Hence, the efficiency ratio is exactly $\frac{3}{5}$.

V. EFFICIENCY WITH MULTI-HOP TRAFFIC FLOWS

In this section, we extend our results to the case when each traffic flow can traverse multiple hops in the network. Let $\lambda^{(s)}$ denote the packet arrival rate of session s and let S denote the

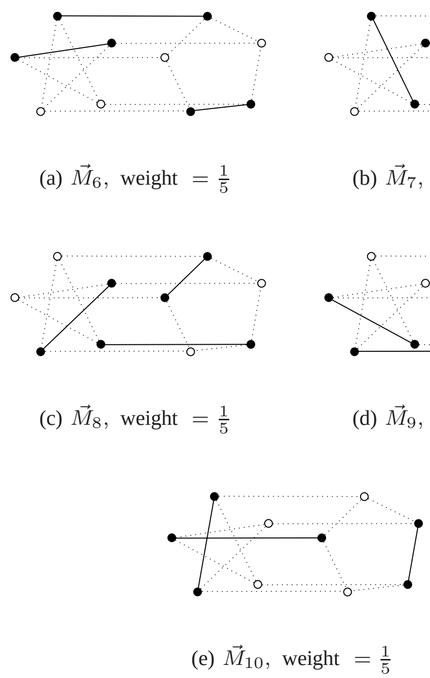

(a) $\vec{M}_6$, weight $= \frac{1}{5}$ (b) $\vec{M}_7$, weight $= \frac{1}{5}$

(c) $\vec{M}_8$, weight $= \frac{1}{5}$ (d) $\vec{M}_9$, weight $= \frac{1}{5}$

(e) $\vec{M}_{10}$, weight $= \frac{1}{5}$

Fig. 6. Maximal matchings for constructing $\vec{\nu}$.

set of sessions s. The optimal capacity region with multi-hop flows is defined as

$$\Omega := \{\vec{\lambda}^{(s)} \mid \vec{\lambda}^{(s)} \preceq \vec{\phi}^{(s)}, \text{ for some } [\textstyle\sum_{s\in S} H_l^{(s)}\phi^{(s)}] \in \Lambda\},$$

where $H_l^{(s)}$ is the routing matrix such that $H_l^{(s)} = 1$ if the path of session s includes link l, and $H_l^{(s)} = 0$ otherwise. Our goal is to modify *GMM* in such a way that it also guarantees the same efficiency ratio γ^* for the multi-hop case.

To this end, we use the idea of regulators proposed in [24], [25]. At each link, there is a separate queue for each session. Let $Q_l^{(s)}(t)$ denote the queue length of session s of link l at the beginning of time slot t, and let $Q_l(t) := \sum_s H_l^{(s)} Q_l^{(s)}(t)$. We add a regulator queue before $Q_l^{(s)}(t)$ as shown in Fig. 7. This additional queue is called a regulator [24], [25] because its service rate is regulated as a function of the mean arrival rate at the source. Hence, the burstiness of the traffic at the down-stream links is reduced specifically. The regulator located at the k-th-hop link ($k \geq 2$), say l, of session s moves accumulated packets to the queue $Q_l^{(s)}(t)$ at the rate of $\lambda^{(s)} + (k-1)\epsilon$ provided that $P_l^{(s)}(t) \geq c_l$, where ϵ is a small positive real number and $P_l^{(s)}(t)$ denotes the number of accumulated packets at the regulator at the beginning of time slot t. Specifically, if $P_l^{(s)}(t) \geq c_l$, the regulator moves c_l packets with probability $\frac{\lambda^{(s)}+(k-1)\epsilon}{c_l}$ at time slot t. We note that, at the first-hop link of each session (where exogenous packets arrive), the packets are directly applied to the service queue. Let l_+^s and l_-^s denote the next hop of link l for session s and the previous hop of link l for session s, respectively. We then have the following equations that govern the evolution of

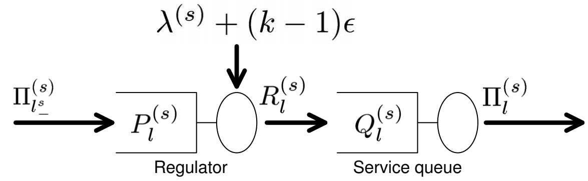

Fig. 7. Description of the network model with regulator at the k-th-hop of session s ($k \geq 2$).

$Q_l^{(s)}(t)$ and $P_l^{(s)}(t)$:

$$Q_l^{(s)}(t+1) = Q_l^{(s)}(t) - \Pi_l^{(s)}(t) + R_l^{(s)}(t), \quad (18)$$

$$P_l^{(s)}(t+1) = P_l^{(s)}(t) - R_l^{(s)}(t) + \Pi_{l_-^s}^{(s)}(t), \quad (19)$$

where $R_l^{(s)}(t)$ denotes the number of packets proceeds from regulator to the queue up to time slot t and $\Pi_l^{(s)}(t)$ denotes the number of packets of session s served at link l up to time slot t. Note that for the first-hop link l_1 of session s, we have $R_{l_1}^{(s)}(t) = A^{(s)}(t)$, where $A^{(s)}(t)$ is the number of packet arrivals of session s up to time slot t. We interpolate the values of these functions to all real-number t by setting $R_l^{(s)}(t) = R_l^{(s)}(\lfloor t \rfloor)$ and $\Pi_l^{(s)}(t) = \Pi_l^{(s)}(\lfloor t \rfloor)$, and by linearly interpolating $Q_l^{(s)}(t)$ and $P_l^{(s)}(t)$ for all t.

Lemma 8: For a network graph with the local-pooling factor σ^*, if the multi-hop traffic rate vector $\vec{\lambda}^{(s)}$ is strictly within $\sigma^*\Omega$, then there exists an $\epsilon > 0$ such that the system is stable under *GMM* with the use of regulators described earlier.

Note that Lemma 8 can be extended to general K-hop interference models since the operations of the regulators do not depend on the underlying interference constraints.

The main idea of the proof is similar to that of [24], [25]. We use the following Lyapunov function, where bold letters indicates a vector.

$$\mathcal{V}(\mathbf{P},\mathbf{Q}) := \xi\mathcal{X}(\mathbf{P},\mathbf{Q}) + \mathcal{Y}(\mathbf{Q}), \quad (20)$$

$$\text{where } \mathcal{X}(\mathbf{P},\mathbf{Q}) := \frac{1}{2}\sum_{l\in E}\sum_{s\in S} H_l^{(s)}(P_{l_+^s}^{(s)} + Q_l^{(s)})^2, \quad (21)$$

$$\mathcal{Y}(\mathbf{Q}) := \max_{l\in E}(Q_l c_l)^2. \quad (22)$$

Note that $\mathcal{Y}(\mathbf{Q})$ is the square of the Lyapunov function for the single-hop case in Section III. The term (21) was introduced in [25] to account for the backlog $P_l^{(s)}$ of the regulators. ξ is a small positive constant to be chosen later. As in (8), we define the fluid limits $q_l^{(s)}(t)$, $p_l^{(s)}(t)$, $\pi_l^{(s)}(t)$, and $r_l^{(s)}(t)$ of $Q_l^{(s)}(t)$, $P_l^{(s)}(t)$, $\Pi_l^{(s)}(t)$, and $R_l^{(s)}(t)$, respectively. The fluid limit version of (20) can also be written as

$$\mathcal{V}(\mathbf{p},\mathbf{q}) = \xi\mathcal{X}(\mathbf{p},\mathbf{q}) + \mathcal{Y}(\mathbf{q}). \quad (23)$$

We can then show that when ϵ is small, the Lyapunov function (23) has a negative drift. We provide the proof in the Appendix.

Proposition 9: For a network graph with the local-pooling factor σ^*, *GMM* along with regulators achieves the efficiency ratio of σ^* under multi-hop traffic load.

Proof: Lemma 8 ensures the system stability with traffic load $\vec{\lambda}^{(s)}$ strictly inside $\sigma^*\Omega$ (i.e., the efficiency ratio is no smaller than σ^*). Hence, it suffices to show that for all $\epsilon > 0$, there exists a scenario that the system is unstable under *GMM* with a traffic rate vector $\vec{\lambda}^{(s)} \in (\sigma^* + \epsilon)\Omega$.

Note that we have shown in Lemma 2 that with single-hop traffic, there exists a traffic pattern that the system is unstable under *GMM*. Since the single-hop traffic scenario can be considered as a special case of the multi-hop traffic scenario, we can use the same technique. We first find a subset L and two vectors $\vec{\mu}, \vec{\nu} \in Co(\mathcal{R}_L)$ such that $\sigma^* \vec{\mu} \succeq \vec{\nu}$. We build a traffic pattern from matchings of $\vec{\nu}$ with $\vec{\lambda}^{(s)} = \vec{\nu} + \epsilon$, as in the proof of Lemma 2, and make each session traverse a single hop. Note that packets are injected at the rate $\lambda^{(s)}$ at each link in L, and are directly applied to the service queue. Then, under *GMM*, the queues of all links in L increase to infinity together as described in the proof of Lemma 2. ■

VI. Conclusion

In this paper, we have provided new analytical results on the achievable performance of *GMM* for a large class of network topologies under the node-exclusive interference model. We derive our results via a topological approach that extends the recently developed notion of local pooling to a more general topological notion called σ-local pooling, and a corresponding notion called local-pooling factor. We show that for a given graph, the efficiency ratio of *GMM* is equal to its local-pooling factor. Thus, we are able to focus on the topological property of graphs to obtain the achievable performance of *GMM*. However, it turns out that estimating the local-pooling factor is non-trivial, and may require high complexity for arbitrary network topologies. Nonetheless, by studying the properties of unstable networks, we can estimate the local-pooling factor of arbitrary network graphs under the node-exclusive interference model and show that the local-pooling factor (and hence the efficiency ratio γ^* of *GMM*) of a graph with maximum node degree d^* is no smaller than $\frac{d^*}{2d^*-1}$. The tightness of the bound is demonstrated through the 6-link cycle and the Star-pentagon topologies, where $d^* = 2$ and $d^* = 3$, respectively. The results are also extended to the more general scenario with multi-hop traffic. We show that when per-flow regulators are used at each hop, *GMM* can also guarantee the same γ^* fraction of the optimal capacity region as in a single-hop traffic scenario.

There remain many interesting open problems in these directions. It would be an interesting avenue for future research to determine whether the bound is tight when $d \geq 4$, and further research on the topological properties of graphs could result in a better estimate of the performance limits. We also expect that different interference models will affect the capacity region of *GMM*. While our results on the relationship between the performance of *GMM* and the local-pooling factor remain the same for a more general class of interference models than the node-exclusive interference model, more work needs to be done to evaluate the local-pooling factor for general interference models [21]. Finally, the authors of [18] show that, if the arrivals satisfy certain randomness property, *GMM* may achieve the full capacity region even if the network graph does not satisfy local pooling. It would be interesting to study whether the results in this paper can be improved under similar assumptions.

Appendix

A. Proof of Lemma 8

We are going to prove that $\frac{D^+}{dt^+}\mathcal{V}(\mathbf{p},\mathbf{q}) < 0$. We first consider the original discrete-time system. We use the convention that all $Q_l^{(s)}(t)$, $P_l^{(s)}(t)$, $R_l^{(s)}(t)$, $\Pi_l^{(s)}(t)$ are equal to zero if $H_l^{(s)} = 0$. From (21), we have

$$\begin{aligned}
&\Delta\mathcal{X}(\mathbf{P},\mathbf{Q}) \\
&\triangleq \tfrac{1}{2}\mathbf{E}\Big[\textstyle\sum_{l\in E}\sum_{s\in S}(P_{l_+^s}^{(s)}(t+1) + Q_l^{(s)}(t+1))^2 \\
&\qquad - \textstyle\sum_{l\in E}\sum_{s\in S}(P_{l_+^s}^{(s)}(t) + Q_l^{(s)}(t))^2 \mid \mathbf{P}(t),\mathbf{Q}(t)\Big] \\
&= \mathbf{E}\Big[\textstyle\sum_l\sum_s(P_{l_+^s}^{(s)}(t) + Q_l^{(s)}(t))\cdot(P_{l_+^s}^{(s)}(t+1) \\
&\qquad + Q_l^{(s)}(t+1) - P_{l_+^s}^{(s)}(t) - Q_l^{(s)}(t)) \mid \mathbf{P}(t),\mathbf{Q}(t)\Big] \\
&\quad + \tfrac{1}{2}\mathbf{E}\Big[\textstyle\sum_l\sum_s(P_{l_+^s}^{(s)}(t+1) + Q_l^{(s)}(t+1) \\
&\qquad - P_{l_+^s}^{(s)}(t) - Q_l^{(s)}(t))^2 \mid \mathbf{P}(t),\mathbf{Q}(t)\Big] \\
&= \mathbf{E}\Big[\textstyle\sum_l\sum_s(P_{l_+^s}^{(s)}(t) + Q_l^{(s)}(t))\cdot(P_{l_+^s}^{(s)}(t+1) \\
&\qquad + Q_l^{(s)}(t+1) - P_{l_+^s}^{(s)}(t) - Q_l^{(s)}(t)) \mid \mathbf{P}(t),\mathbf{Q}(t)\Big] \\
&\quad + \tfrac{1}{2}\mathbf{E}\Big[\textstyle\sum_l\sum_s(R_l^{(s)}(t) - R_{l_+^s}^{(s)}(t))^2 \mid \mathbf{P}(t),\mathbf{Q}(t)\Big].
\end{aligned} \tag{24}$$

Note that $\mathbf{E}[R_l^{(s)}(t)] \leq \lambda^{(s)} + H^*\epsilon$ where H^* is the maximum number of hops of session. Letting $C_1 := |E|\sum_s(\lambda^{(s)} + H^*\epsilon)^2$, we obtain

$$\begin{aligned}
&\Delta\mathcal{X}(\mathbf{P},\mathbf{Q}) \\
&\leq \mathbf{E}\Big[\textstyle\sum_l\sum_s(P_{l_+^s}^{(s)}(t) + Q_l^{(s)}(t))\cdot(P_{l_+^s}^{(s)}(t+1) \\
&\qquad + Q_l^{(s)}(t+1) - P_{l_+^s}^{(s)}(t) - Q_l^{(s)}(t)) \mid \mathbf{P}(t),\mathbf{Q}(t)\Big] + C_1 \\
&\leq \textstyle\sum_l\sum_s(P_{l_+^s}^{(s)}(t) + Q_l^{(s)}(t)) \\
&\qquad \cdot\mathbf{E}\Big[R_l^{(s)}(t) - R_{l_+^s}^{(s)}(t) \mid \mathbf{P}(t),\mathbf{Q}(t)\Big] + C_1 \\
&\leq \textstyle\sum_l\sum_s(P_{l_+^s}^{(s)}(t) + Q_l^{(s)}(t)) \\
&\qquad \cdot(\bar{R}_l^{(s)}\cdot I_{\{P_l^{(s)}(t)\geq c_l\}} - \bar{R}_{l_+^s}^{(s)}\cdot I_{\{P_{l_+^s}^{(s)}(t)\geq c_l\}}) + C_1,
\end{aligned} \tag{25}$$

where $I_{\{\cdot\}}$ is an indicator function, and $\bar{R}_l^{(s)}$ is the average departure rate of the regulator for session s at link l conditioned on $P_l^{(s)}(t) \geq c_l$, except at the first-hop link l_1 of session s, where $\bar{R}_{l_1}^{(s)} = \lambda^{(s)}$. Hence, we obtain

$$\begin{aligned}
&\Delta\mathcal{X}(\mathbf{P},\mathbf{Q}) \\
&\leq \textstyle\sum_l\sum_s(P_{l_+^s}^{(s)}(t) + Q_l^{(s)}(t)) \\
&\qquad \cdot(\bar{R}_l^{(s)} - \bar{R}_{l_+^s}^{(s)}\cdot I_{\{P_{l_+^s}^{(s)}(t)\geq c_l\}}) + C_1 \\
&\leq \textstyle\sum_l\sum_s Q_l^{(s)}(t)c_l \\
&\qquad + \textstyle\sum_l\sum_s P_{l_+^s}^{(s)}(t)(\bar{R}_l^{(s)} - \bar{R}_{l_+^s}^{(s)}\cdot I_{\{P_{l_+^s}^{(s)}(t)\geq c_l\}}) + C_1 \\
&\leq \textstyle\sum_l\sum_s Q_l^{(s)}(t)c_l - \epsilon\sum_l\sum_s P_{l_+^s}^{(s)}(t) + C_2,
\end{aligned} \tag{26}$$

where C_2 is some constant. Note that the last inequality follows from the design of our regulator. Specifically, since $\bar{R}^{(s)}_{l^s_+} = \bar{R}^{(s)}_l + \epsilon$, we have

$$P^{(s)}_{l^s_+}(t)(\bar{R}^{(s)}_l - \bar{R}^{(s)}_{l^s_+} \cdot I_{\{P^{(s)}_{l^s_+}(t) \geq c_l\}}) = \begin{cases} P^{(s)}_{l^s_+}(t)\bar{R}^{(s)}_l < c_l^2, & \text{if } P^{(s)}_{l^s_+}(t) < c_l, \\ -\epsilon P^{(s)}_{l^s_+}(t), & \text{if } P^{(s)}_{l^s_+}(t) \geq c_l, \end{cases} \quad (27)$$

which implies that $P^{(s)}_{l^s_+}(t)(\bar{R}^{(s)}_l - \bar{R}^{(s)}_{l^s_+} \cdot I_{\{P^{(s)}_{l^s_+}(t) \geq c_l\}}) \leq \max_k c_k(c_k + \epsilon) - \epsilon P^{(s)}_{l^s_+}(t)$.

Inequality (26) holds for the original discrete-time system. Now we take the fluid limit as in Section III. Using Theorem 4.1 of [22], we can show that, for almost all sample paths and for all positive sequence $x_n \to \infty$, there exists a subsequence $\{x_{n_j}\}$ such that the following holds uniformly over compact intervals of time t.

$$\begin{aligned} &\tfrac{D^+}{dt^+}\mathcal{X}(\mathbf{p}, \mathbf{q}) = \lim_{x_{n_j} \to \infty} \tfrac{1}{x_{n_j}} \Delta \mathcal{X}(\mathbf{P}(x_{n_j} t), \mathbf{Q}(x_{n_j} t)) \\ &\leq \lim_{x_{n_j} \to \infty} \left(\textstyle\sum_l \sum_s \frac{Q^{(s)}_l(x_{n_j} t) c_l}{x_{n_j}} - \epsilon \sum_l \sum_s \frac{P^{(s)}_{l^s_+}(x_{n_j} t)}{x_{n_j}} + \frac{C_2}{x_{n_j}} \right) \\ &= \textstyle\sum_l \sum_s q^{(s)}_l(t) c_l - \epsilon \sum_l \sum_s p^{(s)}_{l^s_+}(t). \end{aligned} \quad (28)$$

We next derive $\frac{D^+}{dt^+}\mathcal{Y}(\mathbf{q})$. Let $L_0(t) := \{l \in E \mid l = \text{argmax}_k\, q_k(t) c_k\}$ and $L(t) := \{l \in L_0(t) \mid l = \text{argmax}_k \frac{D^+}{dt^+} q_k(t) c_k\}$. Then there exists $\delta > 0$ such that for $(t, t+\delta]$, the links in $L(t)$ have the largest queue-weighted rate. Since $q_l(t)$ is continuous, we have $\frac{D^+}{dt^+}\mathcal{Y}(\mathbf{q}) = 2 q_l(t) c_l^2 \cdot \frac{D^+}{dt^+} q_l(t)$ for any $l \in L(t)$ on the time interval $(t, t+\delta]$. Choose $\epsilon > 0$ sufficiently small so that $[\sum_{s \in S} H^{(s)}_l \lambda^{(s)} + H^* \epsilon]$ is strictly within $\sigma^* \Lambda$. Since *GMM* will always pick a service vector in $Co(\mathcal{R}_{L(t)})$, there exists a link $e \in L(t)$ and a small $\eta > 0$ such that $\frac{D^+}{dt^+} q_e(t) < -\eta$. Hence, we have

$$\tfrac{D^+}{dt^+}\mathcal{Y}(\mathbf{p}, \mathbf{q}) \leq -2\eta q_e(t) c_e^2. \quad (29)$$

Combining (28) and (29), we obtain

$$\begin{aligned} \tfrac{D^+}{dt^+}\mathcal{V}(\mathbf{p}, \mathbf{q}) &= \xi \tfrac{D^+}{dt^+}\mathcal{X}(\mathbf{p}, \mathbf{q}) + \tfrac{D^+}{dt^+}\mathcal{Y}(\mathbf{q}) \\ &\leq \xi \sum_{l \in E} q_l(t) c_l - \epsilon \xi \sum_{l \in E} \sum_{s \in S} p^{(s)}_{l^s_+}(t) - 2\eta q_e(t) c_e^2 \\ &\leq -q_e(t) c_e \cdot (2\eta c_e - \xi |E|) - \epsilon \xi \sum_{l \in E} \sum_{s \in S} p^{(s)}_{l^s_+}(t). \end{aligned} \quad (30)$$

The last inequality holds because $q_l(t) c_l \leq q_e(t) c_e$ for all $l \in E$ and $e \in L(t)$. Hence, if we choose ξ such that $2\eta \cdot \min_k c_k - \xi |E| \geq C_0$ for some constant $C_0 > 0$, then we have that, for all $\mathbf{p}(t) + \mathbf{q}(t) \neq 0$,

$$\tfrac{D^+}{dt^+}\mathcal{V}(\mathbf{p}, \mathbf{q}) \leq -C_0 q_e(t) - \epsilon \xi \sum_{l \in E} \sum_{s \in S} p^{(s)}_{l^s_+}(t) < 0. \quad (31)$$

Since (31) is true for almost every t, the fluid limit model of the system is stable. Then by Theorem 4.2 of [22], the original system is also stable.

REFERENCES

[1] C. Joo, X. Lin, and N. B. Shroff, "Performance Limits of Greedy Maximal Matching in Multi-hop Wireless Networks," in *IEEE CDC*, December 2007.

[2] X. Lin and N. B. Shroff, "The Impact of Imperfect Scheduling on Cross-Layer Congestion Control in Wireless Networks," *IEEE/ACM Trans. Netw.*, vol. 14, no. 2, pp. 302–315, April 2006.

[3] L. Tassiulas and A. Ephremides, "Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximal Throughput in Multihop Radio Networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, December 1992.

[4] C. Joo and N. B. Shroff, "Performance of Random Access Scheduling Schemes in Multi-hop Wireless Networks," in *IEEE INFOCOM*, May 2007.

[5] X. Lin and S. Rasool, "Constant-Time Distributed Scheduling Policies for Ad Hoc Wireless Networks," in *IEEE CDC*, December 2006.

[6] G. Sharma, N. B. Shroff, and R. R. Mazumdar, "Joint Congestion Control and Distributed Scheduling for Throughput Guarantees in Wireless Networks," in *IEEE INFOCOM*, May 2007.

[7] L. Tassiulas, "Linear Complexity Algorithms for Maximum Throughput in Radio Networks and Input Queued Switches," in *IEEE INFOCOM*, April 1998.

[8] S. Ray and S. Sarkar, "Arbitrary Throughput Versus Complexity Tradeoffs in Wireless Networks using Graph Partitioning," in *the Information Theory and Applications Second Workshop*, January 2007.

[9] G. Sharma, C. Joo, and N. B. Shroff, "Distributed Scheduling Schemes for Throughput Guarantees in Wireless Networks," in *the 44th Annual Allerton Conference on Communications, Control, and Computing*, September 2006.

[10] A. Gupta, X. Lin, and R. Srikant, "Low-Complexity Distributed Scheduling Algorithms for Wireless Networks," in *IEEE INFOCOM*, May 2007.

[11] S. Sanghavi, L. Bui, and R. Srikant, "Distributed Link Scheduling with Constant Overhead," in *ACM Sigmetrics*, June 2007.

[12] E. Modiano, D. Shah, and G. Zussman, "Maximizing Throughput in Wireless Networks via Gossiping," *Sigmetrics Performance Evaluation Review*, vol. 34, no. 1, pp. 27–38, 2006.

[13] B. Hajek and G. Sasaki, "Link Scheduling in Polynominal Time," *IEEE Trans. Inf. Theory*, vol. 34, no. 5, September 1988.

[14] S. Sarkar and L. Tassiulas, "End-to-end Bandwidth Guarantees Through Fair Local Spectrum Share in Wireless Ad-hoc Networks," in *IEEE CDC*, December 2003.

[15] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, 1982.

[16] E. Leonardi, M. Mellia, F. Neri, and M. A. Marsan, "On the Stability of Input-Queued Switches with Speed-Up," *IEEE/ACM Trans. Netw.*, vol. 9, no. 1, Feburary 2001.

[17] N. McKeown, "Scheduling Algorithms for Input-Queued Cell Switches," Ph.D. dissertation, University of California at Berkeley, 1995.

[18] 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.

[19] A. Brzezinski, G. Zussman, and E. Modiano, "Enabling Distributed Throughput Maximization in Wireless Mesh Networks: A Partitioning Approach," in *ACM MOBICOM*, 2006, pp. 26–37.

[20] J.-H. Hoepman, "Simple Distributed Weighted Matchings," eprint, October 2004. [Online]. Available: http://arxiv.org/abs/cs/0410047v1

[21] C. Joo, X. Lin, and N. B. Shroff, "Understanding the Capacity Region of the Greedy Maximal Scheduling Algorithm in Multi-hop Wireless Networks," in *IEEE INFOCOM*, April 2008.

[22] J. G. Dai, "On Positive Harris Recurrence of Multiclass Queueing Networks: A Unified Approach via Fluid Limit Models," *Annals of Applied Probability*, vol. 5, no. 1, pp. 49–77, 1995.

[23] C. Joo, X. Lin, and N. B. Shroff, "Greedy Maximal Matching: Performance Limits for Arbitrary Network Graphs Under the Node-exclusive Interference Model," The Ohio State University, Columbus, OH, USA, Tech. Rep., 2009. [Online]. Available: http://www.ece.osu.edu/~cjoo

[24] L. Bui, A. Eryilmaz, R. Srikant, and X. Wu, "Joint Asynchronous Congestion Control and Distributed Scheduling for Multi-Hop Wireless Networks," in *IEEE INFOCOM*, April 2006.

[25] X. Wu, R. Srikant, and J. R. Perkins, "Scheduling Efficiency of Distributed Greedy Scheduling Algorithms in Wireless Networks," *IEEE Trans. Mobile Computing*, vol. 6, no. 6, pp. 595–605, 2007.

Changhee Joo (S'98-M'05) received his Ph.D degree from the school of Electrical and Computer Engineering, Seoul National University in 2005. He was an associated researcher at the Center of Wireless Systems and Applications, Purdue University, from 2005, and joined the Ohio State University in 2007. Currently, he is a research scientist at the Ohio State University. His research interests include multi-hop wireless scheduling, cross-layer design, data aggregation in wireless sensor networks, congestion control, and active queue management. He is a recipient of the IEEE INFOCOM 2008 best paper award, and also a member of the IEEE.

Xiaojun Lin (S'02-M'05) received his B.S. from Zhongshan University, Guangzhou, China, in 1994, and his M.S. and Ph.D. degrees from Purdue University, West Lafayette, Indiana, in 2000 and 2005, respectively. He is currently an Assistant Professor of Electrical and Computer Engineering at Purdue University.

Dr. Lin's research interests are in the analysis, control and optimization of communication networks. He received the 2005 best paper of the year award from *Journal of Communications and Networks* and the IEEE INFOCOM 2008 best paper award. His paper was also one of two runner-up papers for the best-paper award at IEEE INFOCOM 2005. He received the NSF CAREER award in 2007. He was the Workshop co-chair for IEEE GLOBECOM 2007, the Panel co-chair for WICON 2008, and the TPC co-chair for ACM MobiHoc 2009.

Ness B. Shroff (S'91-M'93-SM'01-F'07) received his Ph.D. degree from Columbia University, NY in 1994 and joined Purdue university immediately thereafter as an Assistant Professor. At Purdue, he became Professor of the school of Electrical and Computer Engineering in 2003 and director of CWSA in 2004, a university-wide center on wireless systems and applications. In July 2007, he joined The Ohio State University as the Ohio Eminent Scholar of Networking and Communications, which is considered the highest level of endowed chair at the Ohio State University.

His research interests span the areas of wireless and wireline communication networks. He is especially interested in fundamental problems in the design, performance, control, and security of these networks. Dr. Shroff is a past editor for IEEE/ACM Trans. on Networking and the IEEE Communications Letters. He currently serves on the editorial board of the Computer Networks Journal. He has served on the technical and executive committees of several major conferences and workshops. He was the TPC co-chair of IEEE INFOCOM'03, ACM Mobihoc'08, and general chair of IEEE CCW'99 and WICON'08. Dr. Shroff is a fellow of the IEEE.

He has received numerous awards for his work, including two best paper awards at IEEE INFOCOM (in 2006 and 2008), the flagship conference of the field. He has also received the IEEE IWQoS 2006 best student paper award, the 2005 best paper of the year award for the Journal of Communications and Networking, the 2003 best paper of the year award for Computer Networks, and the NSF CAREER award in 1996 (his IEEE INFOCOM 2005 paper was selected as one of two runner-up papers).