Constant-Time Distributed Scheduling Policies for Ad Hoc Wireless Networks

Xiaojun Lin and Shahzada B. Rasool
School of Electrical and Computer Engineering, Purdue University
West Lafayette, IN 47907, U.S.A.
{linx, srasool}@ecn.purdue.edu.

Abstract— **We propose two new distributed scheduling policies for ad hoc wireless networks that can achieve provable capacity regions. Known scheduling policies that guarantee comparable capacity regions are either centralized or need computation time that increases with the size of the network. In contrast, the unique feature of the proposed distributed scheduling policies is that they are constant-time policies, i.e., the time needed for computing a schedule is independent of the network size. Hence, they can be easily deployed in large networks.**

I. INTRODUCTION

In this paper, we study the link scheduling problem in ad hoc wireless networks. In wireless networks, the radio transmissions at different links can interfere with each other. Hence, in order to achieve the optimal capacity, it is usually more efficient to only use a subset of the radio links at each time [1]. Determining which subset of links should be active at each time becomes the link scheduling problem, which is mainly at the MAC layer in the OSI reference model.

Good scheduling policies are those that can achieve large capacity regions and can be easily computed. Consider a wireless network with L links, and let λ_l be the data rate offered to link l. Let $\vec{\lambda} = [\lambda_1, ..., \lambda_L]$. The *capacity region* under a particular scheduling policy is the set of data-rate vectors $\vec{\lambda}$ that the scheduling policy can support while keeping the queues at all links finite. A scheduling policy is said to be *throughput-optimal* if it can achieve the largest possible capacity region. Known throughput-optimal policies require solving a global optimization problem at each time [2]–[5]. Such scheduling policies are inappropriate for ad hoc networks because the distributive nature of these networks requires simple and decentralized scheduling solutions. Recently, a number of distributed scheduling policies have been proposed in the literature [6]–[10]. Since the capacity region under a distributed policy is typically smaller than the optimal one achieved by the throughput-optimal policy, we define the *efficiency ratio* of such a sub-optimal scheduling policy as the largest number γ such that, given any network topology, for any $\vec{\lambda}$ that can be supported by the throughput-optimal policy, this policy can support $\gamma\vec{\lambda}$. In other words, the scheduling policy with efficiency ratio γ can achieve at least γ fraction of the optimal capacity region. Related works have studied a number of distributed scheduling policies that are shown to achieve provable efficiency ratios. For example,

the maximal matching policy has been shown to achieve an efficiency ratio of no less than $1/2$ under the node-exclusive interference mode [6]. Similar maximal scheduling policies are also studied under the bidirectional-equal-power model and the two-hop interference model [7]–[10], where different bounds on the efficiency ratio are derived.

The problem with these existing distributed scheduling policies, however, is that the time needed to compute a schedule still increases with the size of the network. For example, one of the best known distributed algorithms in graph theory can compute maximal matching on a graph in $O(\log^4 N)$ rounds, where N is the total number of vertices in the graph [11]. Note that we need to be extra careful in interpreting this type of results under the context of ad hoc networks, because these distributed algorithms assume reliable message passing among neighboring nodes in each round of the computation. In ad hoc networks, these messages themselves can interfere with each other. Hence, it is likely that the actual amount of time needed for computing maximal matching in ad hoc networks can be much larger than $O(\log^4 N)$.

In this paper, we propose two new distributed scheduling policies. A unique feature of these new policies is that they are *constant-time* policies, i.e., the time needed to compute a schedule is independent of the size of the network. In fact, our proposed policies only require *one* round of computation. Hence, they are more scalable and easier to implement in large networks. We will provide constant-time distributed scheduling policies for two types of interference models, i.e., the node-exclusive interference model and the two-hop interference model. We will show that our proposed scheduling policies can achieve comparable efficiency ratios as some of the non-constant-time policies in the literature.

II. THE SYSTEM MODEL

Consider a wireless network with N nodes and L links. Each link corresponds to a pair of transmitter node and receiver node. Let $b(l)$ and $e(l)$ denote the transmitter node and the receiver node, respectively, of link l. Two nodes are one-hop neighbors of each other if they are the end-points of a common link. For each node i, let $E(i)$ denote the set of links that connect to the one-hop neighbors of node i, i.e., $E(i)$ is the set of links that node i either acts as a transmitter or as a receiver. Two links are *one-hop* neighbors

of each other if they share a common node. Two links are *two-hop* neighbors of each other if they have a common one-hop neighboring link. For each link l, let $N^1(l)$ denote the set of one-hop neighbors of link l (including link l itself), i.e., $N^1(l) = E(b(l)) \cup E(e(l))$. Further, let $N^2(l)$ denote the set of two-hop neighbors, i.e., $N^2(l) = \bigcup_{k \in N^1(l)} N^1(k)$.

We assume a single-hop traffic model, i.e., each packet only needs to traverse one of the L links and then leave the system. We assume that time is divided into slots of unit length. Let $A_l(t)$ denote the number of packets that arrive at link l at time slot t. We assume that packets are of unit length, and the packet arrival process $A_l(t)$ is stationary and ergodic.

We will study two types of interference models that govern the radio transmission. In both models, we say that two links interfere with each other if they cannot transmit data together. Under the *node-exclusive interference model*, each link l interferes with all of its one-hop neighboring links. Under the *two-hop interference model*, each link l interferes with all of its two-hop neighboring links. In both models, if the above interference constraints are satisfied, an active link l can transfer c_l packets within the time slot. We further assume that the system has carrier-sensing capabilities. In particular, under the one-hop interference model, we assume that all the one-hop neighboring links of link l can sense the transmission at link l. Under the two-hop interference model, we assume that all one-hop neighboring nodes of node i can sense the transmission from node i.

Remark: The node-exclusive interference model can be viewed as a generalization of the bipartite graph model for modeling high-speed packet switches [12], [13]. It has been used in [6], [14], [15] to model wireless networks. While this is a somewhat simplified model, the main results can often be readily generalized to other more complex interference models, e.g., the two-hop interference model. Note also that the latter model is very close to the interference model that IEEE 802.11 DCF (Distributed Coordination Function) deals with [10].

At time slot t, let $\mathcal{M}(t)$ denote the outcome of the scheduling policy, which is defined as the set of non-interfering links that are chosen to be active at time t. Let $D_l(t)$ denote the number of packets that link l can serve at time slot t. Then $D_l(t) = c_l$ if $l \in \mathcal{M}(t)$, and $D_l(t) = 0$ otherwise. Let $Q_l(t)$ denote the number of packets queued at link l at the beginning of time slot t, then the evolution of $Q_l(t)$ is governed by

$$Q_l(t+1) = [Q_l(t) + A_l(t) - D_l(t)]^+, \quad (1)$$

where $[\cdot]^+$ denote the projection to $[0, +\infty)$.

We say that the system is *stable* if the queue lengths at all links remain finite [3], i.e.

$$\lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \mathbf{1}_{\{\sum_{l=1}^{L} Q_l(t) > \eta\}} \to 0, \text{ almost surely as } \eta \to \infty. \quad (2)$$

Let λ_l be the mean packet arrival rate at link l, i.e., $\lambda_l = \mathbf{E}[A_l(t)]$. Let $\vec{\lambda} = [\lambda_1, ... \lambda_L]$. As we defined in the Intro-

duction, the *capacity region* under a particular scheduling policy is the set of $\vec{\lambda}$ such that the system remains stable. The optimal capacity region Ω is the supremum of the capacity regions of all scheduling policies. A scheduling policy is *throughput-optimal* if it can achieve the optimal capacity region Ω. The *efficiency ratio* of a (possibly sub-optimal) scheduling policy is the largest number γ such that the scheduling policy can stablize the system under any load $\vec{\lambda} \in \gamma\Omega$. By definition, a throughput-optimal scheduling policy has an efficiency ratio of 1.

III. A Constant-Time Distributed Scheduling Policy for the Node-Exclusive Interference Model

Two known distributed scheduling policies for the node-exclusive interference mode, i.e., the Maximal Matching Policy [6] and the distributed implementation of the Greedy Maximal Matching Policy [6], [13], have been shown to achieve an efficiency ratio of $1/2$. However, neither of them can compute a schedule in constant time [11], [16]. In this section, we propose a new distributed scheduling policy for the node-exclusive interference model that only needs $O(1)$ time to compute a schedule, and we will show that it achieves an efficiency ratio at least close to $1/3$. The new policy operates as follows:

Constant-Time Distributed Scheduling Policy P:
At each time slot t:
- Each link l computes a probability $p_l(t)$ based on its own queue-length and that of its one-hop neighboring links as follows: $p_l(t) = 0$ if $Q_l(t) = 0$. Otherwise,

$$p_l(t) = \frac{\frac{Q_l(t)}{c_l}}{\max[\sum_{k \in E(b(l))} \frac{Q_k(t)}{c_k}, \sum_{k \in E(e(l))} \frac{Q_k(t)}{c_k}]}. \quad (3)$$

- Each link l attempts transmission with probability $p_l(t)$, and does not attempt transmission with probability $1 - p_l(t)$. For those links that attempt transmission, each of them randomly and independently chooses a backoff time uniformly from $\{0, 1, ..., M-1\}$, where M is a system-wide positive integer constant. We assume that all backoff timers start at the beginning of the time slot. When a link's backoff timer expires, the transmission at the link starts, provided that it has not overheard (i.e., through carrier-sensing) any other transmission from its one-hop neighboring links. Hence, the link l whose backoff timer expires ahead of all of its interfering links will win, and will be able to successfully transfer packets in the time-slot. It is possible that two or more links' backoff timers expire at the same time, in which case collision occurs and none of the interfering links can transfer packets in time-slot t.

The random backoff procedure in the second step of the policy is typical in random access protocols (e.g., IEEE 802.11 and Ethernet) to reduce excessive contention. In practical implementations, the actual backoff time depends both on the constant M and on how long each unit of backoff

time lasts. In practice, due to propagation and processing delays, the length of each unit of backoff time cannot be arbitrarily small. For example, in IEEE 802.11, each unit of backoff time lasts $20\mu s$. Therefore, in order to compute the schedule in constant time, we need to provide an upper bound on M. In Section III-A, we will see how the efficiency ratio of Policy P depends on M.

We also note that Policy P can be viewed as an extension of the Longest-Queue-Driven (LQD) scheduling algorithm from the switching literature [13]. However, there are two key differences: (a) in the switching literature, the network topology is a bipartite graph, while ad hoc network topology is non-bipartite; (b) in the switching literature, the transmitting nodes (i.e., input ports) and receiving node (i.e., output ports) are determined a priori, while in ad hoc networks a node can alternate its role as transmitter or receiver from time-slot to time-slot. The proposed policy P has carefully accounted for these differences through the random backoff phase in the second part of the policy.

A. The Efficiency Ratio of Policy P

We next show that the efficiency ratio of the above policy P is at least close to $1/3$. We start with some definitions.

Definition 1: Let $\vec{x}$ be a component-wise positive vector in $\mathbb{R}^L$ and let Θ be a convex, closed and bounded subset in the positive quadrant of $\mathbb{R}^L$. Assume that Θ contains the intersection of some neighborhood of the origin and the positive quadrant. The *length* of $\vec{x}$ with respect to Θ is defined as

$$||\vec{x}||_\Theta = \frac{1}{\sup\{k|k \geq 0, k\vec{x} \in \Theta\}}.$$

When $\vec{x} \neq 0$, the *normalized* vector $\tilde{\mathbf{x}}$ of $\vec{x}$ with respect to Θ is defined as

$$\tilde{\mathbf{x}} = \frac{\vec{x}}{||\vec{x}||_\Theta}, \quad (4)$$

i.e., $\tilde{\mathbf{x}}$ is the longest vector in Θ that is in the same direction as $\vec{x}$.

The following lemma shows that the above-defined *length* has many properties of a *norm* (except that it is only defined for component-wise positive vectors $\vec{x}$). Its proof is available in our online technical report [17]. Denote $\vec{y} \succeq \vec{x}$ if $\vec{y}$ is component-wise greater than or equal to $\vec{x}$.

Lemma 1:   1) $||\vec{x}|| \geq 0$. Further, $||\vec{x}|| > 0$ if $\vec{x} \neq 0$.
2) $||\tilde{\mathbf{x}}|| = 1$.
3) If $\vec{y} = \alpha\vec{x}$, where $\alpha > 0$, then $||\vec{y}|| = \alpha||\vec{x}||$.
4) If $\vec{y} \succeq \vec{x}$, then $||\vec{y}|| \geq ||\vec{x}||$.
5) $||\vec{x} + \vec{y}|| \leq ||\vec{x}|| + ||\vec{y}||$.

It is well-known that the optimal capacity region Ω under the node-exclusive interference model is upper bounded by [18]:

$$\Psi_0 = \left\{ \vec{\lambda} \ \middle| \ \sum_{l \in E(i)} \frac{\lambda_l}{c_l} \leq 1, \text{ for all nodes } i \right\}. \quad (5)$$

Let $\vec{Q}(t) = [Q_1(t), ..., Q_L(t)]$, where $Q_l(t)$ denote the queue length of link l at the beginning of time slot t. Let $d_l^0(t) = p_l(t)c_l$, and let $\vec{d^0} = [d_1^0, ..., d_L^0]$.

Lemma 2: If $\vec{Q}(t) \neq 0$, then $\vec{d^0}(t) \succeq \tilde{\mathbf{Q}}(t)$, where the normalization is with respect to Ψ_0.

Proof: From (4) and (5), we have $\tilde{\mathbf{Q}}(t) = k_0\vec{Q}(t)$, where k_0 is the largest positive number k that satisfies,

$$\sum_{l \in E(i)} \frac{kQ_l(t)}{c_l} \leq 1 \text{ for all } i.$$

Hence,

$$k_0 = \frac{1}{\max_i \sum_{l \in E(i)} \frac{Q_l(t)}{c_l}}.$$

Using (3), we have $p_l(t) \geq \frac{k_0 Q_l(t)}{c_l}$. The result then follows. $\blacksquare$

Lemma 2 shows that, if links that attempt transmission were to win every time, then the expected amount of service provided by link l at time-slot t is component-wise no less than $\tilde{\mathbf{Q}}(t)$. However, due to the random-backoff procedure in the second part of Policy P, only a few links that attempt transmission will win. We next show that, if a link attempts transmission, the conditional probability that it wins is no less than $\frac{1}{3} - \frac{1}{M}$. In fact, we will prove a more general result as follows. Fix a particular link 0. Label its interfering links as $1, 2, ..., K$. Let x_k denote the probability that the k-th interfering link attempts transmission. Assume that all links follow the random backoff procedure in the second part of Policy P.

Lemma 3: If $\sum_{k=1}^K x_k \leq H$, where $H \geq 0$, then the conditional probability that link 0 wins, conditioned on it attempts transmission, is no less than $\frac{1}{H+1} - \frac{1}{M}$.

Proof: Condition the following derivation on the event that link 0 attempts transmission. Let Y be the random variable that denote the backoff time of link 0, Conditioned on $Y = y$, the probability that link 0 wins is no less than the probability that all K interfering links either do not attempting transmission, or have backoff time greater than y. Note that each interfering link attempts transmission and chooses its backoff time independently. Let S denote the event that link 0 wins. We thus have,

$$\mathbf{P}[S|Y = y] \geq \prod_{k=1}^K \left[\frac{(M-1-y)x_k}{M} + (1 - x_k) \right]$$

$$= \prod_{k=1}^K \left[1 - \frac{y+1}{M}x_k \right].$$

Since Y is uniformly distributed among $\{0, ..., M-1\}$, we have

$$\mathbf{P}[S] = \sum_{y=0}^{M-1} \frac{\mathbf{P}[S|Y=y]}{M}$$

$$\geq \sum_{y=0}^{M-1} \frac{1}{M} \prod_{k=1}^K \left[1 - \frac{y+1}{M}x_k \right].$$

Since $\prod_{k=1}^{K}(1-ux_k)$ is decreasing in u, we have,

$$
\begin{aligned}
\mathbf{P}[S] &\geq \int_{\frac{1}{M}}^{1}\prod_{k=1}^{K}(1-ux_k)\,du \\
&\geq \int_{0}^{1}\prod_{k=1}^{K}(1-ux_k)\,du - \frac{1}{M}. \quad (6)
\end{aligned}
$$

By comparing the derivatives, we can show that

$$
\prod_{k=1}^{K}(1-ux_k) \geq (1-u)^H.
$$

Hence,

$$
\mathbf{P}[S] \geq \int_{0}^{1}(1-u)^H\,du - \frac{1}{M} = \frac{1}{H+1} - \frac{1}{M}.
$$

■

Remark: A special case of Lemma 3 that corresponds to $H = 1$ and $M = \infty$ was shown in Theorem 5 of [13]. Here we have provided a more general result using a much different proof technique.

Under Policy P, we infer from (3) that, for any link l, the attempt probabilities of its one-hop neighboring links must satisfy $\sum_{k\in E(b(l))} p_k(t) \leq 1$ and $\sum_{k\in E(e(l))} p_k(t) \leq 1$. Hence, the sum of the attempt probabilities over its interfering links is no greater than 2. We thus obtain the following corollary to Lemma 3.

Corollary 4: Under Policy P, the conditional probability that link l wins, conditioned on it attempts transmission, is no less than $\frac{1}{3} - \frac{1}{M}$.

Using Lemma 2 and Corollary 4, we thus conclude that the average service provided by link l at time-slot t is no less than the l-th component of $\tilde{\mathbf{Q}}(t)(\frac{1}{3} - \frac{1}{M})$. We can now prove our main result.

Proposition 5: Under Policy P, the network is stable when $\vec{\lambda}$ lies strictly inside the set $(\frac{1}{3} - \frac{1}{M})\Psi_0$.

Proof: We will prove stability by finding a Lyapunov function with negative drift for the fluid model of the system. The fluid model is defined as follows [12], [19]. We first interpolate the values of $Q_l(t)$ to all non-negative real number t by linear interpolation between $\lfloor t \rfloor$ and $\lfloor t \rfloor + 1$ (where $\lfloor t \rfloor$ denote the largest integer no greater than t). We also define the values of $A_l(t)$ and $D_l(t)$ for all real number t by letting $A_l(t) = A_l(\lfloor t \rfloor)$, and $D_l(t) = D_l(\lfloor t \rfloor)$. Then, using the techniques of Theorem 4.1 of [19], we can show that, for almost all sample paths and for any positive sequence $x_n \to \infty$, there exists a subsequence x_{n_j} with $x_{n_j} \to \infty$ such that

$$
\begin{aligned}
&\frac{1}{x_{n_j}}Q_l(x_{n_j}t) \to q_l(t) \text{ for all } l, \\
&\frac{1}{x_{n_j}}\int_{0}^{x_{n_j}t} A_l(s)\,ds \to \lambda_l t \text{ for all } l, \\
&\frac{1}{x_{n_j}}\int_{0}^{x_{n_j}t} D_l(s)\,ds \to \int_{0}^{t} d_l(s)\,ds \text{ for all } l, \quad (7)
\end{aligned}
$$

uniformly over compact intervals. Further, let $\vec{q}(t) = [q_1(t), q_2(t), ...q_L(t)]$ and $\vec{d}(t) = [d_1(t), d_2(t), ..., d_L(t)]$. Using Lemmas 2 and 3, and the techniques of Theorem 4.1 of [19] again, we can show that the limits $\vec{q}(t)$ and $\vec{d}(t)$ satisfy the following set of equations: for all l,

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

$$
q_l(t) \geq 0, \text{ and } \vec{d}(t) \succeq \tilde{\mathbf{q}}(t)(\frac{1}{3} - \frac{1}{M}) \text{ if } \vec{q}(t) \neq 0, \quad (8)
$$

where $\tilde{\mathbf{q}}(t)$ is the normalized vector of $\vec{q}(t)$ with respect to Ψ_0. Any such limit $[\vec{q}(t), \vec{d}(t)]$ is called a *fluid limit* of the system. We say that a fluid limit model of the system is stable if there exists a constant T that depends only on the network topology, the arrival rates λ_l and the active link capacities c_l, such that for *any* fluid limit with $||\vec{q}(0)|| = 1$, we have $||\vec{q}(t)|| = 0$ for all $t \geq T$ [12], [19].

We now use the following Lyapunov function

$$
V(\vec{q}(t)) = ||\vec{q}(t)||
$$

to show that the fluid limit model of the system is stable, where $||\cdot||$ is defined with respect to Ψ_0. Note that for any $\vec{q}(t) \neq 0$, there exists a positive number δt_0, such that for any positive number $\delta t < \delta t_0$, we have $\vec{q}(t) - \vec{d}(t)\delta t \succeq 0$. For such $\vec{q}(t) \neq 0$ and small δt, we have,

$$
\begin{aligned}
&V(\vec{q}(t+\delta t)) \\
\leq\ & ||\vec{q}(t) + \vec{\lambda}\delta t - \vec{d}(t)\delta t|| + o(\delta t) \\
\leq\ & ||\vec{q}(t) - \vec{d}(t)\delta t|| + ||\vec{\lambda}||\delta t + o(\delta t) \\
&\quad \text{(by Part 5 of Lemma 1)} \\
\leq\ & ||\vec{q}(t) - \tilde{\mathbf{q}}(t)(\frac{1}{3} - \frac{1}{M})\delta t|| + ||\vec{\lambda}||\delta t + o(\delta t) \\
&\quad \text{(by (8) and Part 4 of Lemma 1)} \\
=\ & ||\vec{q}(t)|| - (\frac{1}{3} - \frac{1}{M})\delta t + ||\vec{\lambda}||\delta t + o(\delta t) \\
&\quad \text{(by Parts 2 and 3 of Lemma 1).}
\end{aligned}
$$

When $\vec{\lambda}$ lies strictly inside the set $(\frac{1}{3} - \frac{1}{M})\Psi_0$, we have

$$
||\vec{\lambda}|| \leq (\frac{1}{3} - \frac{1}{M}) - \beta,
$$

for some constant $\beta > 0$. Hence, we have,

$$
\frac{D^+ V(\vec{q}(t))}{dt^+} \leq -\beta \text{ if } \vec{q}(t) \neq 0,
$$

where for any function $f(t)$, $\frac{D^+}{dt^+}f(t)$ is defined as $\limsup_{u\downarrow 0}\frac{f(t+u)-f(t)}{u}$. This then shows that the fluid limit model of the system is stable. By Theorem 4.2 of [19], the original system is also stable (i.e., positive Harris recurrent).

■

Remark: For any given $\epsilon > 0$, we can choose the maximum backoff time $M = 1/\epsilon$, which then ensures that the efficiency ratio of Policy P is no less than $1/3 - \epsilon$. Note that for each ϵ, the value of M is independent of the network topology. Hence, we have shown that Policy P only takes constant time and can guarantee an efficiency

ratio close to 1/3 for arbitrary network topologies. As $M \to \infty$, the guaranteed efficiency ratio goes to 1/3. Hence, the difference ϵ is the loss in efficiency due to the constant-time requirement. We also note that the same technique for proving Proposition 5 can be used to establish the result of Theorems 6 and 7 in [13]. Compared with the proofs there, our construction of the Lyapunov function is new and much easier to understand. Alternatively, the result of Proposition 5 may be shown using the Lyapunov function in [20].

IV. A CONSTANT-TIME DISTRIBUTED SCHEDULING POLICY FOR THE TWO-HOP INTERFERENCE MODEL

We next extend the constant-time distributed policy in the previous section to the two-hop interference model. Under the two-hop interference model, the known distributed scheduling policies, i.e., the Maximal Scheduling Policy [7]–[10] and the distributed implementation of the Greedy Scheduling Policy [21], can both guarantee an efficiency ratio of $1/\widehat{N^1}$, where $\widehat{N^1} \triangleq \max_l |N^1(l)|$ is the maximum number of one-hop neighboring links for any link. However, they are again not constant-time policies. We now propose a constant-time distributed scheduling policy Q that can guarantee a comparable efficiency ratio.

Constant-Time Distributed Scheduling Policy Q:

Let W be a positive number between 1 and $\widehat{N^1}$. At each time slot t:

- Each link l computes a probability $p_l(t)$ based on its own queue-length and that of the interfering links as follows: $p_l(t) = 0$ if $Q_l(t) = 0$. Otherwise,

$$p_l(t) = \frac{\frac{Q_l(t)}{c_l}}{\max_{k \in N^1(l)} \sum_{h \in N^1(k)} \frac{Q_h(t)}{c_h}} \times \min\left(1, \frac{W}{\max_{k \in N^1(l)} |N^1(k)|}\right).$$

- Each link l attempts transmission with probability $p_l(t)$, and does not attempt transmission with probability $1 - p_l(t)$. For those links that attempt transmission, each of them randomly chooses a backoff time uniformly from $\{0, 1, ..., M-1\}$. We assume that all backoff timers start at the beginning of the time slot. When the backoff time of a link l expires, the transmitter node $b(l)$ of link l will broadcast an RTS to all of its one-hop neighboring nodes, provided that node $b(l)$ has not overheard any RTS from these one-hop neighboring nodes. Once the receiver node $e(l)$ correctly receive the RTS, it will then respond with a CTS broadcasted to all of its neighboring nodes. Through this RTS-CTS procedure, the link l that sends out an RTS before any of its two-hop neighboring links will win. This link L can then transfer packets at the rate of c_l during the rest of the time slot. It is possible that two or more links in a two-hop neighborhood send out RTS together, in which case collision occurs and none of the interfering links can transfer data in time-slot t.
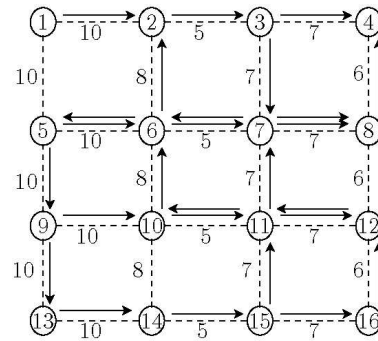

Fig. 1. Network topology

We can use similar technique as in Section III to show the following main result. Readers can refer to [17] for more details.

Proposition 6: Under Policy Q, the network is stable when $\vec{\lambda}$ lies strictly inside the set $\frac{W}{\widehat{N^1}}\left(\frac{1}{1+W} - \frac{1}{M}\right)\Psi_0'$.

Remark: For any fixed W, by choosing $M \geq \frac{1}{(1+W)(2+W)}$, the efficiency ratio of Policy Q is at least $\frac{W}{(2+W)\widehat{N^1}}$. For each W, the value of M is independent of the network topology. Hence, we have shown that Policy Q only takes constant time. As $W \to \widehat{N^1}$ and $M \to \infty$, the guaranteed efficiency ratio goes to $1/(1 + \widehat{N^1})$. The difference $\left(\frac{1}{1+\widehat{N^1}} - \frac{W}{(2+W)\widehat{N^1}}\right)$ is the loss of capacity due to the constant-time requirement.

V. SIMULATION RESULTS

We have simulated the proposed scheduling policies using the network topology in Fig. 1. There are 16 nodes (represented by circles) and 24 links (represented by dashed lines). The capacity is labeled next to each link. The flows are represented by arrows. We simulate single-hop flows, and we let the rate of each flow be λ. Note that although the rates of the flows are the same, the link capacities and the flows have been chosen to avoid uniform patterns.

We first simulate Policy P for the node-exclusive interference model. In Fig. 2, we plot the mean total queue backlog summed over all links of the network, as the offered load λ increases. When λ approaches a certain limit, the average total backlog will increase to infinity. This limit can then be viewed as the boundary of the capacity region. We have plotted the curves for Policy P with maximum backoff window $M = 1$, $M = 10$, and $M = 20$. We can see that the performance of the scheduling policy is much worse when $M = 1$. Hence, the random backoff procedure is essential. However, once M is above a reasonable number, the performance will be virtually the same (as we can see for $M = 10$ and $M = 20$). We have also plotted the performance of the Maximal Matching (MM) policy and the Greedy Maximal Matching (GMM) policy [6]. Although the efficiency ratio that can be guaranteed in Proposition 5 is slightly worse than that of MM, the simulation results indicate that their actual performance are roughly the same.
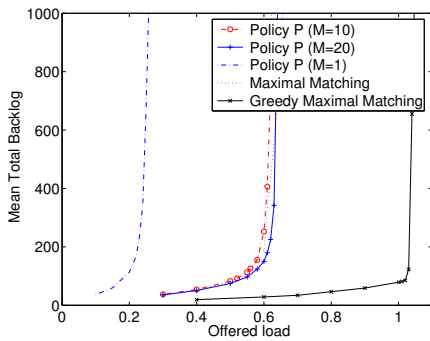
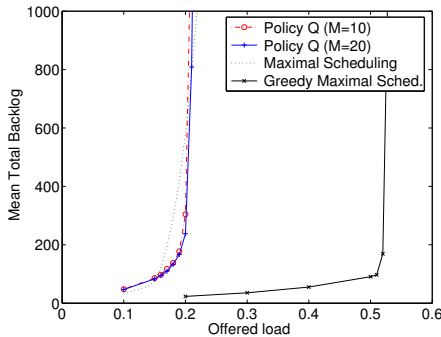Fig. 2. Performance comparison under the node-exclusive interference model

Fig. 3. Performance comparison under the two-hop interference model

We next simulate Policy Q for the two-hop interference mode and plot the results in Fig. 3. Again, we observe that the performance of policy Q changes little when the maximum backoff window changes from $M = 10$ to $M = 20$. Further, the performance is also comparable to the maximal scheduling policy [7], [10].

VI. Conclusion

In this paper, we have proposed two new distributed scheduling policies for ad hoc wireless networks. The unique feature of these new distributed scheduling policies is that they are constant-time policies, i.e., the time needed for computing a schedule is independent of the network size. Hence, they can be easily deployed in large networks. We have shown that both scheduling policies can guarantee efficiency ratios comparable to other known distributed scheduling policies in the literature.

The analysis in the paper could be extended to multi-hop networks and to take into account the delay in exchanging local queue-length information (see the discussions in [17]). We believe that these results offer new insights for the design of simple and efficient scheduling policies for ad hoc networks. We also observe in both Fig. 2 and Fig. 3 that the performance of Greedy Maximal Scheduling policies is typically better than that of Policies P and Q, even though their provable efficiency ratios are comparable. For future work, it would be interesting to study whether one can develop constant-time distributed scheduling policies that

achieve comparable performance as these Greedy Maximal Scheduling policies.

References

[1] P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 388–404, March 2000.

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

[3] M. J. Neely, E. Modiano, and C. E. Rohrs, "Dynamic Power Allocation and Routing for Time Varying Wireless Networks," in *Proceedings of IEEE INFOCOM*, San Francisco, April 2003.

[4] A. Eryilmaz, R. Srikant, and J. Perkins, "Stable Scheduling Policies for Fading Wireless Channels," *IEEE/ACM Transactions on Networking*, vol. 13, no. 2, pp. 411–424, April 2005.

[5] R. L. Cruz and A. V. Santhanam, "Optimal Routing, Link Scheduling and Power Control in Multi-hop Wireless Networks," in *Proceedings of IEEE INFOCOM*, San Francisco, April 2003.

[6] X. Lin and N. B. Shroff, "The Impact of Imperfect Scheduling on Cross-Layer Rate Control in Multihop Wireless Networks," in *Proceedings of IEEE INFOCOM*, Miami, FL, March 2005.

[7] P. Chaporkar, K. Kar, and S. Sarkar, "Throughput Guarantees in Maximal Scheduling in Wireless Networks," in *Proceedings of 43rd Annual Allerton Conference on Communication, Control and Computing*, Monticello, IL, September 2005.

[8] ——, "Achieving Queue Length Stability Through Maximal Scheduling in Wireless Networks," in *Proceedings of Information Theory and Applications Inaugural Workshop*, University of California, San Diego, February 2006.

[9] X. Wu, R. Srikant, and J. R. Perkins, "Queue-Length Stability of Maximal Greedy Schedules in Wireless Network," in *Proceedings of Information Theory and Applications Inaugural Workshop*, University of California, San Diego, February 2006.

[10] X. Wu and R. Srikant, "Bounds on the Capacity Region of Multihop Wireless Networks Under Distributed Greedy Scheduling," in *Proceedings of IEEE INFOCOM*, Barcelona, Spain, April 2006.

[11] M. Hanckowiak, M. Karonski, and A. Panconesi, "On the Distributed Complexity of Computing Maximal Matching," *SIAM Journal of Discrete Mathematics*, vol. 15, no. 1, pp. 41–57, 2001.

[12] J. G. Dai and B. Prabhakar, "The Throughput of Data Switches with and without Speedup," in *Proceedings of IEEE INFOCOM*, Tel Aviv, Israel, March 2000, pp. 556–564.

[13] E. Leonardi, M. Mellia, F. Neri, and M. A. Marsan, "On the Stability of Input-Queued Switches with Speed-Up," *IEEE/ACM Transactions on Networking*, vol. 9, no. 1, pp. 104–118, February 2001.

[14] S. Sarkar and L. Tassiulas, "End-to-end Bandwidth Guarantees Through Fair Local Spectrum Share in Wireless Ad-hoc Networks," in *Proceedings of the IEEE Conference on Decision and Control*, Maui, Hawaii, December 2003.

[15] Y. Yi and S. Shakkottai, "Hop-by-hop Congestion Control over a Wireless Multi-hop Network," in *Proceedings of IEEE INFOCOM*, Hong Kong, March 2004.

[16] D. E. Drake and S. Hougardy, "A Linear-Time Approximation Algorithm for Weighted Matchings in Graph," *ACM Transactions on Algorithms*, vol. 1, no. 1, pp. 107–122, July 2005.

[17] X. Lin and S. Rasool, "Constant-Time Distributed Scheduling Policies for Ad Hoc Wireless Networks," *Technical Report, Purdue Uniersity, available at http://min.ecn.purdue.edu/~linx/papers.html*, 2006.

[18] B. Hajek and G. Sasaki, "Link Scheduling in Polynomial Time," *IEEE Transactions on Information Theory*, vol. 34, no. 5, pp. 910–917, September 1988.

[19] J. G. Dai, "On Positive Harris Recurrence of Multiclass Queueing Networks: A Unified Approach via Fluid Limit Models," *Annals of Applied Probability*, vol. 5, pp. 49–77, 1995.

[20] A. Eryilmaz, R. Srikant, and J. R. Perkins, "Stable Scheduling Policies for Broadcast Channels," in *ISIT 2002*, Lausanne, Switzerland, June-July 2002.

[21] G. Sharma, R. R. Mazumdar, and N. B. Shroff, "On the Complexity of Scheduling in Wireless Networks," in *ACM Mobicom*, Los Angeles, CA, September 2006.