

Capacity of Asynchronous Random-Access Scheduling in Wireless Networks

Deepti Chafekar*, Dave Levin[†], V. S. Anil Kumar[‡], Madhav V. Marathe,[‡] Srinivasan Parthasarathy[§] and Aravind Srinivasan[¶]

*Department of Computer Science, Virginia Tech, Blacksburg, VA 24061, Email: chafekar@vt.edu

[†]Department of Computer Science, University of Maryland, College Park, MD 20742, Email: dml@cs.umd.edu

[‡]Department of Computer Science and Virginia Bioinformatics Institute, Virginia Tech, Blacksburg, VA 24061
Email:{akumar, mmarathe}@vbi.vt.edu

[§]IBM T.J. Watson Research Center, Hawthorne, NY 10532, Email: spartha@us.ibm.com

[¶]Dept. of Computer Science and Institute for Advanced Computer Studies,
University of Maryland, College Park, MD 20742, Email: srin@cs.umd.edu

Abstract—We study the throughput capacity of wireless networks which employ (asynchronous) random-access scheduling as opposed to deterministic scheduling. The central question we answer is: how should we set the channel-access probability for each link in the network so that the network operates close to its optimal throughput capacity? We design simple and distributed channel-access strategies for random-access networks which are provably competitive with respect to the optimal scheduling strategy, which is deterministic, centralized, and computationally infeasible. We show that the competitiveness of our strategies are nearly the best achievable via random-access scheduling, thus establishing fundamental limits on the performance of random-access. A notable outcome of our work is that random access compares well with deterministic scheduling when link transmission durations differ by small factors, and much worse otherwise. The distinguishing aspects of our work include modeling and rigorous analysis of asynchronous communication, asymmetry in link transmission durations, and hidden terminals under arbitrary link-conflict based wireless interference models.

I. INTRODUCTION

Two fundamental problems in communication networks are as follows: (i) What is the capacity of the network under a given throughput objective? Specifically, given a collection of communicating source-destination pairs (or connections) in the network, what is the maximum throughput¹ which the network can deliver to the connections? (ii) How can we allocate individual connection throughputs, route the data from the source of each connection to its corresponding destination, and schedule the transmissions at each link in order to operate the network close to its capacity? These problems are particularly complicated in wireless networks due to interference constraints, which prevent proximate links in the network from being active simultaneously. These constraints make it computationally infeasible to determine the optimal routing and scheduling strategies in a wireless network.

Many recent results in the literature, such as by Kumar *et al.* [13], Jain *et al.* [8], Alicherry *et al.* [1], Kodialam and Nandagopal [11], Sharma *et al.* [17] and Wang *et al.* [18] have addressed the problems of wireless capacity estimation

and throughput maximization under various models of interference; they have presented joint routing and link scheduling strategies which are guaranteed to utilize the network *close* to its optimal capacity.

Unfortunately, the scheduling disciplines developed in the above approaches are *deterministic*, and share some common disadvantages. For instance, these scheduling algorithms assume that the network is perfectly time-synchronized; all links agree on the index of the current slot and when the next slot begins. Further, nodes in deterministic protocols exchange (possibly large) lists of time slots during which they transmit, so that other interfering nodes may choose their transmit times carefully and avoid conflicts. As a consequence, even when the throughput demand changes for a *single* link, the network may need to recompute the schedule for *every* link and disseminate the new schedule. In contrast, in a *random-access network*, we need to specify only a single parameter for each node (or link): its channel access probability. Once we instantiate the channel access probabilities at each node, a *distributed* stochastic process which is local to each node governs the use of the common wireless medium, and automatically determines the per-node throughputs.

Motivated by the pervasiveness of random-access scheduling protocols such as 802.11, we consider in this work the problems of capacity estimation and throughput maximization in the context of *asynchronous random-access wireless networks*. Specifically, consider a link-throughput vector, whose components specify the throughput achieved by each link in the network. ***The central contribution of this work is the solution to the following problem: given a link-throughput vector which can be achieved by an optimal scheduling strategy, what should be the channel access probability for each link so that the resultant random-access throughput vector is (component-wise) close to the optimal link-throughput vector?*** It is well known that by solving this fundamental link-scheduling problem and combining it with standard network-flow linear programming formulations, we can solve the more general end-to-end routing, scheduling, and utility maximization problem (see for example [13], [8],

¹or more generally, throughput based utility

[1], [11], [17], [18]). Our focus is therefore on link-scheduling strategies for asynchronous random-access networks that perform *provably close* to optimal scheduling.

A salient aspect of our work is that the access strategies we develop are oblivious to the throughput vector. Not only does each link achieve its desired throughput via an appropriate choice of channel access probability, but it does so without any knowledge of the throughput demands of its interfering links.

We now introduce several definitions to describe our technical contributions in greater detail. A *link-rate vector* \vec{f} is a vector whose components specify the steady-state throughputs we need to support on each link of the network. The *rate-region of a scheduling algorithm* denotes the set of all link-rate vectors which can be supported by the algorithm. A scheduling protocol \mathcal{A} is α -*competitive*² w.r.t. a protocol \mathcal{B} if, whenever \mathcal{B} can achieve a link-rate vector \vec{f} , \mathcal{A} can achieve at least an α -fraction of \vec{f} .³ The *network interference degree*, denoted by Δ , is defined as the maximum number of links which interfere with some specific link ℓ , but are mutually interference-free amongst themselves; this is a key parameter which will frequently appear in our performance guarantees. Further, a key contribution of a majority of the deterministic protocols mentioned above [13], [1], [11], [17] is the design of *deterministic* scheduling strategies whose competitiveness is $\frac{1}{\Delta}$ w.r.t. optimal scheduling; their differences lie mainly in the interference models they consider.

A. Our contributions

(a) Strategies for synchronous random-access. To demonstrate the key elements of our approach, we begin our analysis in Section IV with a *synchronous* model of random-access. A simple non-linear programming (NLP) formulation exists in this scenario, which precisely characterizes the achievable rate-region of synchronous random-access. However, this rate-region is *non-convex*, which introduces significant obstacles in using existing convex-programming techniques for throughput optimization in random-access networks. We design a synchronous random access strategy, and prove that it is $\frac{1}{e\Delta}$ -competitive w.r.t. optimal scheduling.⁴

(b) A first strategy for asynchronous random-access. We study the asynchronous model of random-access in Section V by incorporating the effect of hidden terminals and non-uniform transmission durations. We design an asynchronous random-access strategy and show that it is $\frac{1}{e\Delta(\gamma+1)}$ -competitive w.r.t. optimal scheduling; here, γ is the maximum, taken over all links ℓ in the network and over all hidden links ℓ' that interfere with ℓ , of the ratio of the transmission durations of ℓ and ℓ' .

(c) Fundamental limits of random-access. We show that the competitiveness of *any* random-access strategy is bounded above by (essentially) $\frac{1}{\Delta\gamma}$. That is, no matter how clever

we are in the choice of per-link channel access probabilities, we cannot improve upon the (essentially) $\frac{1}{\Delta\gamma}$ -competitive factor of asynchronous random-access scheduling. This factor therefore represents a *fundamental performance limit* on the competitiveness of random-access (and is not an artifact of our strategies or analysis).

(d) Capacity of random-access. We develop novel necessary and sufficient conditions that are linear, and which characterize the achievable rate-region of an asynchronous random-access network with a given topology and link transmission durations. The gap between our necessary and sufficient conditions is $e\Delta$. In the context of end-to-end utility maximization, plugging the sufficient conditions into the network-flow formulation (as in [13], [8], [1], [11], [17], [18]) immediately yields the end-to-end throughputs, the routing, as well as a random-access link scheduling strategy which approximates the optimal random-access capacity to within a factor of $e\Delta$.

We observe that for many *geometric* models of interference studied in recent literature, it is known that the parameter Δ is upper-bounded by a fixed constant in any network; our techniques allow us to obtain *constant factor performance guarantees* under these models. For instance, for the uniform Tx-model [13], the Tx-Rx model of interference with parameters 1, 2, and 2.5 [1], the K -hop interference model on geometric graphs [17], and the node-exclusive interference model [17], [9], [14], our *synchronous* random-access strategy has a competitive ratio of $5e$, $4e$, $8e$, $12e$, $49e$, and $2e$ respectively, and our *asynchronous* random-access scheduling strategy has a competitive ratio of $O(\gamma)$ w.r.t. optimal scheduling.

A notable consequence of our results **(b)** and **(c)** is that random-access is generally more competitive when packet sizes for each link are chosen in proportion to the capacity of the link (so that link transmission durations become uniform).

We validate our theoretical insights using the CPLEX solver and NS-2 simulations of the 802.11 random-access protocol. Our main observations are: (i) the feasible rate region of random access is indeed non-convex even in simple settings, and hence unlikely to be computationally feasible. This necessitates the use of approximate but provably-good scheduling strategies and capacity estimation techniques such as the ones developed in this work; (ii) the stability conditions we develop accurately model the dynamics of asynchronous random-access in multi-hop topologies; (iii) random-access is generally more competitive as packet sizes are chosen in proportion to link capacities; this is especially evident in high-traffic regimes where the network experiences significant interference; lastly (iv) our new capacity estimation techniques yield much tighter bounds on the capacity of an asynchronous random-access network than existing formulations.

II. RELATED WORK

Capacity of wireless networks: Estimating the capacity of a wireless network and developing communication strategies which operate the network close to its capacity is a complex task due to wireless interference. Recent works which address this problem include those of Jain *et al.* [8], Kodialam and

²Alternatively, competitiveness (or) competitive-ratio of \mathcal{A} w.r.t. \mathcal{B} .

³As per convention, α -fraction of a vector \vec{f} implies that we are multiplying each component of \vec{f} with the scalar α .

⁴ e denotes the base of natural logarithms throughout this paper.

Nandagopal [10], [11], Kumar *et al.* [13], Alicherry, Bhatia, and Li [1], Balakrishnan *et al.* [2], Sharma, Mazumdar, and Shroff [17] and Wang *et al.* [18]. With one exception [8], these results are constructive; not only do they estimate the network capacity, but they also present routing and scheduling algorithms (and in the multi-channel cases [11], [1], channel assignment algorithms as well) which can jointly achieve the capacity. However, the scheduling disciplines underlying all of these results are *deterministic*, and their results are not applicable to random-access networks.

Network capacity with random-access scheduling: The recent results of Lin and Rasool [14], and Joo and Shroff [9] are most similar in spirit to our work, and deal with the performance of *synchronous* random-access for the special-cases of node-exclusive and two-hop interference models. In the node-exclusive model, links interfere with each other only if they share a common end-point; Lin and Rasool [14] design a random-access scheduling strategy based on the notion of periodic contention-frames, which is guaranteed to achieve a (nearly) $\frac{1}{3}$ -factor of the optimal capacity region of the network. Joo and Shroff [9] propose an improved strategy for the node-exclusive interference model which achieves a (nearly) $\frac{1}{2}$ -factor of the optimal capacity region. Both of these works also study the two-hop interference model and propose synchronous random-access schemes which are guaranteed to achieve a (nearly) $\frac{1}{\Delta+1}$ -factor of the optimal capacity, where Δ is the network interference degree.

Our work departs from the results of [14], [9] in the following significant ways. Foremost, our analysis of synchronous random-access is only a building block for our analysis of *asynchronous* random-access. Synchronization is hard if not impossible to achieve in practice (without special hardware such as GPS); often, a central purpose of randomness in scheduling is to overcome a lack of network synchronization. Our work handles two complex challenges that *do not* arise in the analysis of synchronous wireless networks—the effect of hidden terminals and asymmetric link transmission durations—and initiates the capacity analysis of asynchronous random-access. Further, we study the performance of random-access under *arbitrary* link-conflict-based models of wireless interference. For synchronized networks, we show that random-access is $\frac{1}{e\Delta}$ -competitive w.r.t. optimal scheduling under *any* link-conflict based interference model; for asynchronous networks, we show $\frac{1}{e\Delta(\gamma+1)}$ -competitive strategies.

Single-hop random access networks: Random-access protocols in general and 802.11 in particular have received extensive treatment for the special case of a single-hop network (i.e., every node can communicate with every other node, and every link interferes with every other link) [3], [4], [5], [7], [12], [15], [16]. The crucial distinction between these studies and our work is that we need to address the asymmetry that is inherent in multi-hop networks. In a single-hop network all nodes perceive the wireless channel to be in the same state: occupied or idle. Such symmetry does not hold in a multi-hop network which introduces significantly more difficult challenges.

III. BACKGROUND

A. Network Model

We model the random-access wireless network as a directed graph $G = (V, E)$. Each link $\ell \in E$ has a fixed capacity $c(\ell)$ which denotes the maximum bit-rate at which data can be transmitted across ℓ . Link ℓ employs a fixed packet size $M(\ell)$ for its transmission, and a single transmission by ℓ lasts for $T_{xmit}(\ell) = \frac{M(\ell)}{c(\ell)}$ units of time. A transmission across link ℓ will be successful if and only if it is not lost due to channel errors, and if no other interfering link transmits simultaneously. We model channel errors by a parameter $\rho(\ell)$ which denotes the probability that a transmission on link ℓ will *not* encounter any channel errors (i.e., $1 - \rho(\ell)$ is the probability that a channel error occurs during a transmission on ℓ). We define an interference set $I(\ell)$ for link ℓ ; this consists of the set of all links which interfere with link ℓ . A transmission by ℓ is interference-free only if it does not overlap in time with another transmission by a link in its interference set $I(\ell)$. In our general (non-geometric) interference model, the interference set $I(\ell)$ can be specified to be an *arbitrary* subset of links $E \setminus \{\ell\}$. To model hidden links, we partition the interference set $I(\ell)$ into two subsets: $hidden(\ell) \subseteq I(\ell)$, and $exposed(\ell) = I(\ell) \setminus hidden(\ell)$. If link ℓ is currently not involved in a transmission, then it can sense any ongoing transmission by the links in the set $exposed(\ell)$ and vice-versa. However, ℓ *cannot* sense any transmission by the links in the set $hidden(\ell)$ and vice-versa. Thus, we could have scenarios where a link $\ell' \in hidden(\ell)$ starts its transmission in the middle of an ongoing transmission by ℓ ; this will result in the failure of the transmission by ℓ .

B. Asynchronous random-access model of scheduling

Our asynchronous random-access model of scheduling is inspired by the basic access mechanism (or the ad-hoc mode operation) of the 802.11 protocol. Each link perceives time to be slotted and divided into contiguous chunks of length T_{id} . Link ℓ attempts to access the channel during time chunk i as follows. Let time chunk j be such that $j > i$. If chunk j is not part of an ongoing transmission by any link in the set $\{\ell\} \cup exposed(\ell)$, then ℓ *independently* initiates a transmission of the packet during chunk j with probability $\tau(\ell)$. With the remaining probability $1 - \tau(\ell)$, ℓ skips the chunk j . Once ℓ initiates a transmission, it occupies the channel for $T_{xmit}(\ell)$ duration of time. This process is repeated until the packet is successfully transmitted. In general, the channel access probability $\tau(\ell)$ is a function of the rate demanded by ℓ , the capacity $c(\ell)$, and other parameters.⁵

From the perspective of link ℓ , the channel consists of two types of periods: (1) an idle period consisting of a *sequence* of idle chunks (or idle slots) each of length T_{id} , followed by (2) a busy period or a single *busy* slot of length T_{busy} , during which the channel is either occupied by ℓ , some link(s) in $exposed(\ell)$, or a combination thereof. Note that a link *freezes*

⁵Our results can be stated in terms of node-based random-access process as well; we use a link-based process for ease of exposition.




Figure 1. Timing diagram illustrating asynchronous random-access. Links ℓ_2 and ℓ_3 both interfere with ℓ_1 ; ℓ_2 is exposed while ℓ_3 is hidden. Small white rectangles represent idle slots, shaded rectangles transmission slots, and striped rectangles times at which the access process is frozen. The interfering transmission on hidden link ℓ_3 starts in the middle of an ongoing transmission on ℓ_1 ; in contrast, if the exposed link ℓ_2 detects any transmission on ℓ_1 , it freezes its attempt process and vice-versa. A transmission on ℓ_2 can interfere with a transmission on ℓ_1 only if they start during the same chunk.

its attempt process when it is not transmitting but the channel is busy due to some link in $exposed(\ell)$. Also, while each idle slot is of a fixed duration T_{id} , a busy slot has a variable duration that depends upon which link(s) are transmitting and how their transmissions overlap. For ease of analysis, we will assume that the T_{xmit} values are integral multiples of T_{id} . We present in Figure 1 a sample illustration of the random-access process, with a special emphasis on the role of exposed and hidden links.

C. Optimal vs. random-access scheduling

We now present a simple necessary condition which must be satisfied by any throughput vector that is achievable through optimal scheduling. This will be useful in the competitive analysis of our random-access strategies. Suppose we are given a link-throughput vector (or link-rate vector) $\vec{f} = \langle f(\ell_1), \dots, f(\ell_m) \rangle$ where $f(\ell_i)$ specifies the steady state rate we need to support on link ℓ_i . Recall that $\rho(\ell)$ denotes the probability of a transmission *not* failing due to channel errors, and $c(\ell)$ denotes the capacity of link ℓ . Let $x(\ell) = \frac{f(\ell)}{\rho(\ell)c(\ell)}$ denote link ℓ 's *utilization*. A link's utilization is the fraction of the time it needs to be *active without interference* in any schedule which achieves the link-rate vector \vec{f} .

Clearly, not every link-rate vector can be feasibly scheduled and additional constraints are necessary to ensure that it can be scheduled. Recall that $I(\ell)$ denotes the set of links which interfere with link ℓ ; further, recall that Δ is the maximum number of links in any set $I(\ell)$ that are mutually interference-free. It is well-known (and can be easily verified) that every link-rate vector \vec{f} that can be achieved through optimal scheduling must satisfy the following **necessary** condition:

$$\forall \ell : \frac{f(\ell)}{\rho(\ell)c(\ell)} + \sum_{\ell' \in I(\ell)} \frac{f(\ell')}{\rho(\ell')c(\ell')} \leq \Delta \quad (1)$$

Eqn. (1) states that for any link ℓ , the total utilization of ℓ and its interfering links can never exceed Δ . Specifically, if a rate vector violates this constraint, then no scheduling algorithm can achieve this rate vector. We will exploit this

fact to establish the competitiveness of our random-access scheduling strategies w.r.t. optimal scheduling.

Unlike optimal scheduling, which is allowed to choose the transmit slots for each link carefully in a deterministic conflict-free manner, the only parameter which can be tuned for each link in random-access scheduling is its channel-access probability. We are interested in analyzing the set of all link-rate vectors achievable through all possible settings of the channel-access probabilities of the network links.

For ease of presentation, we make the following simplifying assumptions regarding network stability. If a link is scheduled for transmission at a slot when its packet-queue is empty, then the link will transmit a dummy packet of suitable size. A link rate vector is said to be *stable* (or *feasible*), if a strategy exists for scheduling the links such that, at each link, the expected volume of data which is successfully transmitted per slot equals or exceeds the expected volume of data which arrives exogenously at this link. Under certain mild conditions on the traffic arrival process, the stability proofs in this paper can be extended for multi-hop routing and can be stated rigorously in a queueing theoretic setting, using the Lyapunov function framework as in the work of Georgiadis *et al.* [6].

IV. SYNCHRONOUS RANDOM-ACCESS

We now demonstrate the key elements of our approach with synchronous random-access, in which time is divided into equally sized slots of unit length, a transmission on any link spans one slot, and slots across links are synchronized.

A. Synchronous random-access rate-region

We start by characterizing the exact conditions under which a given link rate vector \vec{h} can be achieved through synchronous random-access. Recall that $\rho(\ell)$ is the probability of link ℓ being free from channel errors during a transmission. Let $\eta(\ell)$ denote the probability that no collision occurs at link ℓ (due to interference) during a transmission, and $\tau(\ell)$ the probability with which link ℓ attempts transmission during a slot. We now present the NLP which precisely characterizes the feasible link-rate region of synchronous random-access.

$$\forall \ell \in E : h(\ell) = \tau(\ell)\rho(\ell)\eta(\ell)c(\ell) \quad (2)$$

$$\forall \ell \in E : \eta(\ell) = \prod_{\ell' \in I(\ell)} (1 - \tau(\ell')) \quad (3)$$

Eqn. (2) is the stability condition which states that the rate at which data arrives at link ℓ is equal to the rate at which data is successfully transmitted out of link ℓ . The l.h.s. of Eqn. (2) denotes the expected volume of data (in bits) that enters link ℓ during each (unit length) time slot. The r.h.s. of Eqn. (2) denotes the expected volume of data transmitted on ℓ during each (unit length) time slot. This incorporates $\tau(\ell)$: the probability of a transmission being attempted on ℓ during a slot, $\rho(\ell)\eta(\ell)$: the probability of the transmission succeeding, and $c(\ell)$: the number of bits transmitted during the (unit length) slot if the transmission succeeds. Since transmission events across links are independent, we can express $\eta(\ell)$ as a function of the interfering links' τ 's. Each link ℓ' in the interfering set $I(\ell)$ chooses to transmit during a time slot, independently

at random, with probability $\tau(\ell')$. The transmission at link ℓ is interference-free if and only if none of the links in $I(\ell)$ transmit in the same slot as ℓ . Eqn. (3) captures this. A link-rate vector \vec{h} can be achieved through synchronous random-access if and only if it is a feasible solution to (2) and (3). In general, the set of feasible rate-vectors characterized by (2) and (3) is non-convex. This is a significant obstacle to directly plugging-in known convex-programming-based optimization techniques for network optimization, and necessitates the use of approximate techniques developed in this work.

B. Competitive Scheduling

Here, we develop our synchronous random-access scheduling strategy and prove that it is $\frac{1}{e\Delta}$ -competitive w.r.t. optimal scheduling. Suppose we are given a link-rate vector \vec{f} which is guaranteed to be achievable through optimal scheduling (hence, it satisfies Eqn. (1)). In our random-access strategy, we assign access probabilities to each link ℓ as follows:

$$\forall \ell \in E : \tau(\ell) = 1 - e^{-\frac{f(\ell)}{\Delta \cdot \rho(\ell) \cdot c(\ell)}} \quad (4)$$

Theorem 1: Let \vec{f} be a link-rate vector which satisfies Eqn. (1). If we assign the channel access probabilities for each link using Eqn. (4), then synchronous random-access achieves a link-rate vector \vec{h} such that each component of \vec{h} is at least $\frac{1}{e\Delta}$ times the corresponding component in \vec{f} .

Proof: Let us first compute $\eta(\ell)$ the probability of collision-free transmission for link ℓ as follows:

$$\begin{aligned} \eta(\ell) &= \prod_{\ell' \in I(\ell)} (1 - \tau(\ell')) = e^{-\sum_{\ell' \in I(\ell)} \frac{f(\ell')}{\Delta \cdot \rho(\ell') \cdot c(\ell')}} \\ &\geq e^{-\frac{f(\ell)}{\Delta \cdot \rho(\ell) \cdot c(\ell)} - 1} \end{aligned} \quad (5)$$

The last inequality follows since \vec{f} satisfies (1). Let \vec{h} be the link-rate vector achieved by random-access. We have: $\forall \ell \in E : h(\ell) = \rho(\ell)c(\ell)\tau(\ell)\eta(\ell)$ {from (2)} $\geq \rho(\ell)c(\ell) \cdot (1 - e^{-\frac{f(\ell)}{\Delta \cdot c(\ell) \cdot \rho(\ell)}}) \cdot e^{-\frac{f(\ell)}{\Delta \cdot \rho(\ell) \cdot c(\ell)} - 1}$ {from (5)} $= \rho(\ell)c(\ell) \cdot (e^{\frac{f(\ell)}{\Delta \cdot \rho(\ell) \cdot c(\ell)} - 1} - e^{-1}) \geq \rho(\ell)c(\ell) \cdot \left(\frac{1 + \frac{f(\ell)}{\Delta \cdot \rho(\ell) \cdot c(\ell)}}{e} - \frac{1}{e}\right)$ {since $\forall x : e^x \geq 1 + x$ } $= \frac{f(\ell)}{e\Delta}$. This completes the proof of the theorem. ■

V. ASYNCHRONOUS RANDOM-ACCESS

We begin our study of the asynchronous random-access model by extending the stability condition (2) to asynchronous networks in Section V-A. In general, it seems impossible to completely specify the stability condition for asynchronous networks, since it is not easy to obtain a closed form expression for the success probability $\eta(\ell)$. However, our specification will play a crucial role in the analysis of our scheduling strategies in Section V-B.

A. Stability condition

Let the expected length of a time slot for a link ℓ in steady-state be $\overline{T(\ell)}$. Recall that $T_{xmit}(\ell)$ is the length of a transmit slot for link ℓ . If the link-flow vector supported by the network is \vec{h} , it follows that the expected number of bits generated per time slot for link ℓ is equal to $h(\ell)\overline{T(\ell)}$. Recall that

$(1 - \rho(\ell))$ is the probability of a channel error occurring during a transmission across link ℓ , and $(1 - \eta(\ell))$ is the probability of interference occurring at link ℓ during a transmission. The expected number of bits successfully transmitted over link ℓ per time slot is $c(\ell)T_{xmit}(\ell)$ times the probability of a successful transmission on ℓ during that transmit time slot; i.e., $\tau(\ell) \cdot \rho(\ell)\eta(\ell) \cdot c(\ell)T_{xmit}(\ell)$. This product differs from (2) in that it incorporates $c(\ell)T_{xmit}(\ell)$: the number of bits transmitted over this link-dependent period of time.

We are now ready to state our partial stability constraint, which states that the expected number of bits generated at link ℓ is equal to the expected number of bits which are successfully transmitted on ℓ during each slot:

$$\forall \ell \in E : h(\ell)\overline{T(\ell)} = \tau(\ell)\rho(\ell)\eta(\ell)c(\ell)T_{xmit}(\ell) \quad (6)$$

Eqn. (6) is a partial stability condition since we are not aware of a closed form expression for $\eta(\ell)$ in terms of the τ values. However, this partial specification suffices for our analysis.

B. Competitive scheduling

We now prove our main result that asynchronous random-access is $\frac{1}{e\Delta(\gamma+1)}$ -competitive w.r.t. optimal scheduling. Recall that $T_{xmit}(\ell)$ denotes the length of a transmit slot for link ℓ , and T_{id} denotes the length of an idle slot. Recall also that γ denotes the maximum ratio between the T_{xmit} values of a link ℓ and a hidden link ℓ' in its interference neighborhood: i.e., $\gamma = \max_{\ell} \max_{\ell' \in hidden(\ell)} \frac{T_{xmit}(\ell)}{T_{xmit}(\ell')}$ (note that $\gamma \geq 1$). Suppose we are given a link-rate vector \vec{f} which is guaranteed to be achievable through optimal scheduling (it satisfies Eqn. (1)). Let $\epsilon = \frac{1}{\gamma+1}$. In our asynchronous random-access strategy, we assign access probabilities to each link ℓ as follows:

$$\forall \ell \in E : \tau(\ell) = 1 - e^{-\epsilon \cdot \frac{f(\ell)}{\Delta \cdot c(\ell) \cdot \rho(\ell)} \cdot \frac{T_{id}}{T_{xmit}(\ell)}} \quad (7)$$

The following theorem shows that this choice of access probabilities achieves the rate-vector $\frac{\vec{f}}{(\gamma+1) \cdot e\Delta}$.

Theorem 2: Let \vec{f} be a link-rate vector which satisfies Eqn. (1). If we assign the channel access probabilities for each link using Eqn. (7), then asynchronous random-access achieves a link-rate vector \vec{h} such that each component of \vec{h} is at least $\frac{1}{(\gamma+1) \cdot e\Delta}$ times the corresponding component in \vec{f} .

Proof: Consider a fixed link ℓ . We will now obtain a lower bound on the rate achieved by link ℓ through the strategy in Eqn. (7). Let $\alpha(\ell) \in [0, 1]$ be defined through Eqn. (8).

$$\frac{f(\ell)}{\Delta \cdot c(\ell) \rho(\ell)} + \sum_{\ell' \in exposed(\ell)} \frac{f(\ell')}{\Delta \cdot c(\ell') \rho(\ell')} = \alpha(\ell) \quad (8)$$

$$\sum_{\ell' \in hidden(\ell)} \frac{f(\ell')}{\Delta \cdot c(\ell') \rho(\ell')} \leq 1 - \alpha(\ell) \quad (9)$$

Consider a transmission by link ℓ which starts at time chunk k . This transmission spans the chunks $k, k+1, \dots, k + \frac{T_{xmit}(\ell)}{T_{id}} - 1$, and will be successful if the following conditions hold: (i) No link $\ell' \in hidden(\ell)$ starts a transmission during the chunks $k - \frac{T_{xmit}(\ell')}{T_{id}} + 1, k - \frac{T_{xmit}(\ell')}{T_{id}} + 2, \dots, k, k+1, \dots, k + \frac{T_{xmit}(\ell)}{T_{id}} - 1$, and (ii) No link $\ell' \in exposed(\ell)$ starts a transmission during chunk k .

The probability of a link ℓ' starting a transmission at any specific chunk is at most $\tau(\ell')$, even when conditioned on any other transmission event. Let $\psi \doteq \max_{\ell} \frac{T_{id}}{T_{xmit}(\ell)}$. We have,

$$\eta(\ell) \geq \frac{\prod_{\ell' \in \text{hidden}(\ell)} (1 - \tau(\ell'))}{\prod_{\ell' \in \text{exposed}(\ell)} (1 - \tau(\ell))} \left(\frac{T_{xmit}(\ell) + T_{xmit}(\ell')}{T_{id}} - 1 \right) \times \prod_{\ell' \in \text{exposed}(\ell)} \left(e^{-\frac{\epsilon \cdot T_{id}}{T_{xmit}(\ell')} \cdot \frac{f(\ell')}{\Delta \cdot c(\ell') \rho(\ell')}} \right) \left(\frac{T_{xmit}(\ell) + T_{xmit}(\ell')}{T_{id}} - 1 \right) \times \prod_{\ell' \in \text{exposed}(\ell)} \left(e^{-\frac{\epsilon T_{id}}{T_{xmit}(\ell')} \cdot \frac{f(\ell')}{\Delta \cdot c(\ell') \rho(\ell')}} \right). \quad \text{This along with Eqns. (8) and (9) implies:}$$

$$\eta(\ell) \geq e^{-\epsilon \cdot (1 - \alpha(\ell)) \cdot (1 + \gamma)} \cdot e^{-\epsilon \psi \left(\alpha(\ell) - \frac{f(\ell)}{\Delta \cdot c(\ell) \rho(\ell)} \right)} \quad (10)$$

Link ℓ perceives any time chunk as being in one of the following two states: (i) the chunk is idle, or else (ii) the chunk is occupied by either ℓ or some link $\ell' \in \text{exposed}(\ell)$. In general, the channel alternates between an *occupied period* or an *occupied slot* (consisting of contiguous occupied chunks) and an *idle period* (consisting of contiguous idle chunks or idle slots). Let $\kappa_{idle}(\ell)$ denote the expected number of chunks in an idle period of link ℓ ; let $\kappa_{occ}(\ell)$ denote the expected number of chunks in an occupied period of link ℓ . Let $p_{idle}(\ell)$ denote the steady state probability of a chunk being idle for link ℓ . Therefore, we have $p_{idle}(\ell) = \frac{\kappa_{idle}(\ell)}{\kappa_{idle}(\ell) + \kappa_{occ}(\ell)}$. Recall that $\overline{T}(\ell)$ denotes the expected length of a time slot for link ℓ . Hence,

$$\overline{T}(\ell) \leq \frac{T_{id} \cdot (\kappa_{idle}(\ell) + \kappa_{occ}(\ell))}{\kappa_{idle}(\ell) + 1} \leq \frac{T_{id}}{p_{idle}(\ell)} \quad (11)$$

The intuition behind the first inequality is that, in steady state, link ℓ perceives $\kappa_{idle}(\ell)$ slots of length T_{id} which are idle, followed by a single occupied slot of length $\kappa_{occ}(\ell)T_{id}$ (in expectation).

Observe that link ℓ perceives any fixed chunk k to be idle if and only if ℓ does not start a transmission during the slots $k - \frac{T_{xmit}(\ell)}{T_{id}} + 1, \dots, k$ and no link $\ell' \in \text{exposed}(\ell)$ starts a transmission during the slots $k - \frac{T_{xmit}(\ell')}{T_{id}} + 1, \dots, k$. Hence, we have: $p_{idle}(\ell) \geq (1 - \tau(\ell))^{\frac{T_{xmit}(\ell)}{T_{id}}} \cdot \prod_{\ell' \in \text{exposed}(\ell)} (1 - \tau(\ell'))^{\frac{T_{xmit}(\ell')}{T_{id}}} \geq e^{-\epsilon \alpha(\ell)}$. Combining this with Eqn. (11) yields:

$$\overline{T}(\ell) \leq T_{id} e^{\epsilon \alpha(\ell)} \quad (12)$$

From Eqns. (7) and (10), we have:

$$\begin{aligned} \tau(\ell)\eta(\ell) &\geq (1 - e^{-\epsilon \cdot \frac{f(\ell)}{\Delta \cdot c(\ell) \rho(\ell)} \cdot \frac{T_{id}}{T_{xmit}(\ell)}}) \times \\ &\quad e^{-\epsilon \cdot (1 - \alpha(\ell)) \cdot (1 + \gamma)} \cdot e^{-\epsilon \psi \left(\alpha(\ell) - \frac{f(\ell)}{\Delta \cdot c(\ell) \rho(\ell)} \right)} \\ &= e^{-\epsilon \cdot (1 - \alpha(\ell)) \cdot (1 + \gamma) - \epsilon \psi \alpha(\ell)} \times \\ &\quad e^{\frac{\epsilon \psi f(\ell)}{\Delta \cdot c(\ell) \rho(\ell)}} \times (1 - e^{-\epsilon \cdot \frac{f(\ell)}{\Delta \cdot c(\ell) \rho(\ell)} \cdot \frac{T_{id}}{T_{xmit}(\ell)}}) \\ &\geq e^{\alpha(\ell) - 1 - \epsilon \psi \alpha(\ell)} \times \left(e^{\epsilon \cdot \frac{f(\ell)}{\Delta \cdot c(\ell) \rho(\ell)} \cdot \frac{T_{id}}{T_{xmit}(\ell)}} - 1 \right) \\ &\quad \left\{ \text{since } \epsilon = \frac{1}{\gamma + 1}, \text{ and } \forall \ell: \psi \geq \frac{T_{id}}{T_{xmit}(\ell)} \right\} \\ &\geq e^{\alpha(\ell) - 1 - \epsilon \psi \alpha(\ell)} \cdot \epsilon \cdot \frac{f(\ell)}{\Delta \cdot c(\ell) \rho(\ell)} \cdot \frac{T_{id}}{T_{xmit}(\ell)} \quad (13) \end{aligned}$$

{since $\forall x: e^x \geq 1 + x$ }

We are now ready to bound the rate $h(\ell)$ achieved by link ℓ in steady-state. We have, $h(\ell) = \frac{\tau(\ell)\eta(\ell)\rho(\ell)T_{xmit}(\ell)c(\ell)}{T(\ell)}$ {from Eqn. (6)} $\geq \frac{e^{\alpha(\ell) - 1 - \epsilon \psi \alpha(\ell)} \cdot \epsilon \cdot \frac{f(\ell)}{\Delta \cdot c(\ell) \rho(\ell)} \cdot \frac{T_{id}}{T_{xmit}(\ell)} \cdot \rho(\ell) T_{xmit}(\ell) c(\ell)}{T(\ell)}$ $\geq e^{-\epsilon \alpha(\ell)}$ {from Eqn. (13)} $\geq \epsilon \frac{f(\ell)}{\Delta} e^{\alpha(\ell) - 1 - \epsilon \psi \alpha(\ell)}$ $\cdot e^{-\epsilon \alpha(\ell)}$ {from Eqn. (12)} $= \frac{f(\ell)}{(\gamma + 1) \cdot e \Delta}$ $\cdot e^{\alpha(\ell) - \epsilon \alpha(\ell) - \epsilon \psi \alpha(\ell)}$ {since $\epsilon = \frac{1}{\gamma + 1}$ } $\geq \frac{f(\ell)}{(\gamma + 1) \cdot e \Delta}$ {as $\psi \sim 0$, and $\epsilon \leq 1$ }. Here the assumption $\psi \sim 0$ simply reflects the fact that, in practice, the length of the idle slot is negligible in comparison with the length of transmission durations. This assumption is by no means necessary for our analysis: relaxing it affects our competitive ratio only by a minor factor of $1 + \psi$. This completes the proof of the theorem. ■

VI. FUNDAMENTAL LIMITS OF RANDOM-ACCESS

Theorem 2 states our choice of channel-access probabilities ensure that the competitive ratio of asynchronous random-access w.r.t. optimal scheduling is $\frac{1}{e \Delta (\gamma + 1)}$. We now claim that this ratio is essentially the best achievable and cannot be significantly improved in general, even by a more clever choice of channel access probabilities. Given a network \mathcal{N} with a link set \mathcal{L} , define $\gamma(\mathcal{N}) = \max_{\ell \in \mathcal{L}} \max_{\ell' \in \text{hidden}(\ell)} \frac{T_{xmit}(\ell)}{T_{xmit}(\ell')}$; further, let $\Delta(\mathcal{N})$ denote the interference degree of the specific network \mathcal{N} . We show:

Theorem 3: Consider the family of networks containing all networks \mathcal{N} such that $\gamma(\mathcal{N}) \geq \gamma^*$ and $\Delta(\mathcal{N}) \geq \Delta^*$, where $\gamma^* \geq 1$ is any fixed constant and $\Delta^* \geq 1$ is any fixed integer. The competitiveness of asynchronous random-access w.r.t. optimal scheduling is upper bounded by $\Phi(\Delta^*, \gamma^*) = \frac{2(1 + \ln(\Delta^* \gamma^*))}{\Delta^* \gamma^*}$ for this family of networks.

Proof: Let \mathcal{N} be a network with $\Delta^* + 1$ links $\ell_0, \ell_1, \ell_2, \dots, \ell_{\Delta^*}$; we will name ℓ_0 as the *long* link and the rest of the links as the *short* links. The short links are mutually interference-free; but each of them is a hidden link for the long link and vice-versa. Let all the short links have a transmit duration T_{xmit} ; let $T_{xmit}(\ell_0) = \gamma^* T_{xmit}$. Let $T_{xmit} = \zeta T_{id}$, where $\zeta \geq 1$ is a fixed integer. Let all the links have unit capacities, and let the channel be error-free (hence losses are entirely due to interference). It is easy to see that the link-rate vector $\vec{f} = \langle f(\ell_0), f(\ell_1), \dots, f(\ell_{\Delta^*}) \rangle = \langle \frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2} \rangle$ is achievable in this network through optimal scheduling.

Assume that there exist channel access probabilities which can support the rate vector $\Phi(\Delta^*, \gamma^*) \vec{f}$. The expected fraction of the time link any short link ℓ_i is transmitting on the channel is $\frac{\tau(\ell_i) T_{xmit}}{\tau(\ell_i) T_{xmit} + (1 - \tau(\ell_i)) T_{id}}$. The link capacity is one unit, and ℓ_i achieves a rate of at least $\frac{\Phi(\Delta^*, \gamma^*)}{2}$; hence, $\frac{\tau(\ell_i) T_{xmit}}{\tau(\ell_i) T_{xmit} + (1 - \tau(\ell_i)) T_{id}} \geq \frac{\Phi(\Delta^*, \gamma^*)}{2}$. Combining this with the fact that $T_{id} \leq T_{xmit}$ and by rearranging the terms, we have, for all short links ℓ_i :

$$\tau(\ell_i) \geq \frac{\Phi(\Delta^*, \gamma^*)}{2\zeta} \quad (14)$$

As in the proof of Theorem 2, consider time being divided into chunks of length T_{id} . A transmission by the long link ℓ_0 occupies $\gamma^*\zeta$ contiguous chunks; this transmission will be successful only if no earlier transmission by any short link ℓ_i overlaps with these chunks *and* ℓ_i does not start a transmission in any of these chunks. It follows that $\eta(\ell_0)$, the probability of a successful transmission by ℓ_0 is:

$$\begin{aligned} \prod_{i \in \{1, 2, \dots, \Delta^*\}} (1 - \tau(\ell_i))^{\gamma^*\zeta} &\leq \prod_{i \in \{1, 2, \dots, \Delta^*\}} e^{-\gamma^*\zeta\tau(\ell_i)} \leq e^{-(1 + \ln \Delta^*)\gamma^*\zeta} \leq \frac{1}{e^{\Delta^*\gamma^*}} \end{aligned} \quad (15)$$

It is easy to verify that when the link capacity of ℓ_0 is one unit, the rate achievable by ℓ_0 is upper bounded by $\eta(\ell_0)$; this value is at most $\frac{1}{e^{\Delta^*\gamma^*}}$ by Eqn. (15). Further, it is easy to check that for all $\gamma^* \geq 1$, $\frac{1}{e^{\Delta^*\gamma^*}} < \frac{\Phi(\Delta^*, \gamma^*)}{2}$. Hence, if the rate achieved by ℓ_1 is at least $\frac{\Phi(\Delta^*, \gamma^*)}{2}$, then the rate achieved by ℓ_0 is strictly less than $\frac{\Phi(\Delta^*, \gamma^*)}{2}$. This contradicts our assumption and completes the proof of the theorem. ■

VII. CAPACITY OF RANDOM ACCESS

We now address the central cross-layer optimization problem of maximizing end-to-end connection throughput. The crux of the end-to-end optimization problem lies in characterizing the achievable link-rate region of a network efficiently; here, we present our novel *linear* characterization of the link rate-region that is achievable in a given network through asynchronous random-access. Due to lack of space, we present these characterizations in Theorems 4 and 5, without proof.

Theorem 4: Every link-rate vector \vec{f} which can be stably scheduled through asynchronous random-access satisfies the following linear condition:

$$\forall \ell \in E : \frac{f(\ell)}{\rho(\ell)c(\ell)} + \sum_{\ell' \in \text{exposed}(\ell)} \frac{f(\ell')}{\rho(\ell')c(\ell')} + \sum_{\ell' \in \text{hidden}(\ell)} \frac{f(\ell')}{\rho(\ell')c(\ell')} \cdot \frac{T_{xmit}(\ell') + T_{xmit}(\ell) - T_{id}}{T_{xmit}(\ell')} \leq \Delta \quad (16)$$

Theorem 5: Every link-rate vector \vec{f} which satisfies the following linear condition can be stably scheduled through asynchronous random-access:

$$\forall \ell \in E : \frac{f(\ell)}{\rho(\ell)c(\ell)} + \sum_{\ell' \in \text{exposed}(\ell)} \frac{f(\ell')}{\rho(\ell')c(\ell')} + \sum_{\ell' \in \text{hidden}(\ell)} \frac{f(\ell')}{\rho(\ell')c(\ell')} \cdot \frac{T_{xmit}(\ell') + T_{xmit}(\ell) - T_{id}}{T_{xmit}(\ell')} \leq \frac{1}{e} \quad (17)$$

We can now combine our linear constraints in Eqn. (17) with the standard network flow formulations (as in [13], [8], [1], [11], [17], [18]) to compute the end-to-end capacity of an asynchronous random-access wireless network efficiently.

VIII. SIMULATION RESULTS

We present our experimental results from NS-2 simulations of 802.11 as well as capacity computations using the CPLEX linear programming solver.

A. Validating the stability condition

We begin our experimental evaluation by validating the stability condition (6) in a large multi-hop wireless network. We created a random 200 node network spread over a 1500×1500 square grid. We fixed the capacity of each link in the network to either 24Mbps or 6Mbps with equal probability in order to provide rate diversity. We picked 16 links at random and created a random link-demand vector \vec{f} ; at time t , we injected data into an active link ℓ at the rate $g(t)f(\ell)$. Here, $g(t)$ is a network-wide parameter that increases monotonically with time, while $f(\ell)$ is one of the 16 components in the demand vector \vec{f} . We fixed packet sizes at 512 Bytes and ran the simulation for 120 secs. During each second, we measured $\overline{T}(\ell)$: the average length of a time slot for link ℓ , $\tau(\ell)$: the channel access probability for link ℓ , and $\eta(\ell)$: the probability of interference-free transmission, on all the 16 active links ℓ .

Figure 2(a) shows the observed and predicted channel access probability for a representative link as a function of time. The predictions make use of the observed $\eta(\ell)$ and $\overline{T}(\ell)$ values and plug them into Eqn. (6) to estimate $\tau(\ell)$. *The predicted channel access probability closely mirrors the observed channel access probability until ~ 92 secs. This is the point at which this link saturates, and hence the stability Eqn. (6) is no longer applicable.* Figure 2(b) plots the CDF of the relative-error of our prediction compared to our observations for *all active links* in the network. The three curves correspond to three snapshots in time when the load on the network is 25%, 50% and 75% respectively of the saturation load. *Over 90% of our predictions have significantly less than 2% relative error and thus confirms that Eqn. (6) indeed captures the dynamics of a stable (yet-to-be saturated) random-access network.*

B. Non-convexity

Figure 2(c) illustrates the non-convexity of the random-access rate-region. In this simulation, we used a four node linear network with two active link flows. The links were node-disjoint; the two senders interfered with each other's receiver but were hidden from each other. We set the channel to be error-free, with packet losses due solely to interference. We fixed the bit-rates of both the links to 6 Mbps and varied the packet sizes from 500 Bytes to 2000 Bytes (each plot is for a fixed packet size). We fixed the data arrival rate for the first flow (x coordinate) and observed the maximum throughput achievable by the second flow (y coordinate) *without* affecting the throughput of the first flow. *For any fixed packet size, the achievable random-access rate-region is the area below its corresponding curve. This is distinctly non-convex, and the extent of non-convexity increases as we increase the packet size and hence, the effect of interference.*

C. Competitive Analysis

We now study how the relative value of link transmission durations affects the competitiveness of random-access. We start with a simple setting in Figure 3(a). Here, we created a network with two hidden interfering links ℓ_1 and ℓ_2 , with 6Mbps and 24Mbps capacities respectively. We fixed the

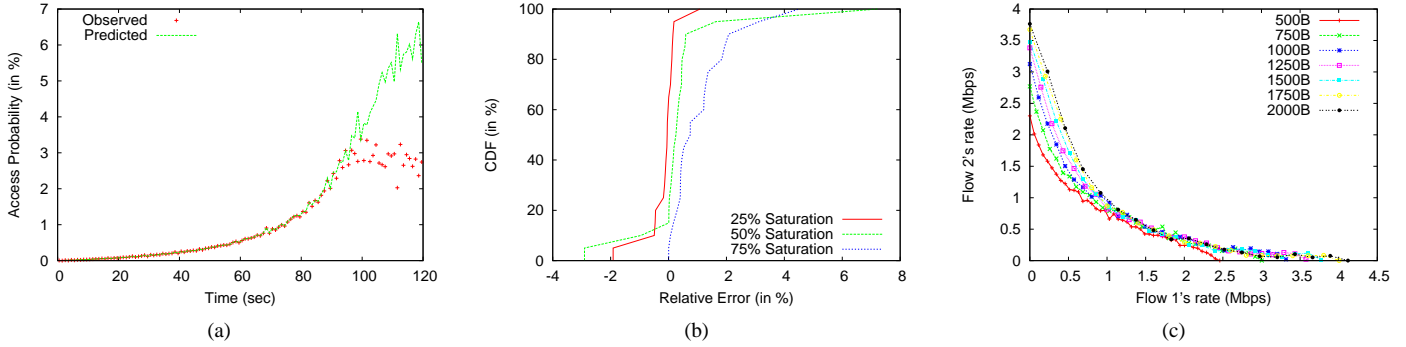


Figure 2. Validating the stability condition and the non-convexity of random-access rate-region: (a) The predicted channel access probability for a representative link closely matches its observed access probability (b) The CDF of the relative error between predicted and observed channel access probability for all the links: more than 90% of predictions have a relative error of less than 2% (c) The random-access rate-region is distinctly non-convex.




Figure 3. Competitive analysis of uniform and proportional packet size policies: (a) Proportional packet size policy has a significantly larger rate-region than uniform packet size policy (b) Proportional packet size policy generally delivers greater throughput fractions than uniform packet size policy *despite* its increased per-packet overhead (c) The competitiveness of uniform packet size size rapidly drops compared to proportional packet size policy as the number of flows, and hence the extent of interference increases in the network.

packet size on link ℓ_1 to 500 Bytes and varied ℓ_2 's packet size from 500 to 2000 Bytes. The two extremes correspond to the uniform packet size policy (500B), and the proportional policy (2000B), where the packet sizes are chosen in proportion to the link capacities. We fix the data arrival rate for the first flow (x coordinate) and observe the maximum throughput achievable by the second flow (y coordinate) *without* affecting the throughput of the first flow. *The proportional packet size policy leads to a significantly larger achievable rate-region than the uniform packet size policy.*

We extend competitive analysis of uniform and proportional policies for the large 200-node multi-hop wireless network considered in Section VIII-A. In the uniform packet size policy, we set the packet sizes for all links at 2000 Bytes; in the proportional policy, we set the packet sizes for the slower links at 500 Bytes, which *increases* per-packet transmission overheads. In each run of the simulation, we picked k random links and a random link-demand vector f ; we varied the number of links k between 2 and 20. At time t , we injected data at an active link ℓ at the rate $g(t)f(\ell)$, $g(t)$ is a monotonically increasing parameter, similar to the validation experiments. We computed the throughput fraction g_1 for the uniform policy, which is the value of $g(t)$ when the network became saturated; we also computed the throughput fraction g_2

for the proportional policy. The value $\frac{g_1}{g_2}$ is the competitiveness of uniform policy compared to proportional policy for this run of the experiment.

Figure 3(b) presents the CDFs of the throughput fractions delivered by the uniform packet size policy (2000B, 2000B), and the proportional packet size policy (500B, 2000B). Clearly, *the proportional packet size policy delivers significantly larger throughput fractions than the uniform policy.* We show the average competitiveness of the uniform policy w.r.t. the proportional policy in Figure 3(c). *Random-access with proportional packet sizes is generally more competitive than random-access with uniform packet sizes; this is especially true when the number of flows in the network increases and the network experiences substantial interference.*

D. Capacity Estimation

As above, we partitioned the network links into fast links and slow links. We assigned the fast links $54Mbps$ capacity, and the slow links a capacity of $\frac{54Mbps}{\gamma}$, where $\gamma \geq 1$ is a variable simulation parameter. We selected two end-to-end connections with arbitrarily chosen source and destination nodes, and computed the total throughput that could be delivered by the network using the CPLEX solver. We used two distinct linear programs (LPs): LP1 uses the mathematical programming formulation of [13], [1], [11], [17] as-is,




Figure 4. Existing capacity formulations (LP1) do not account for relative values of link transmission durations and could significantly overestimate the achievable throughput. Our new capacity formulation (LP2) accounts for the disparity in link transmission durations.

without incorporating the capacity loss due to random-access scheduling. LP2 uses the constraints developed in Section VII.

Figure 4 presents the computed capacity as a function of γ . Since the LP formulation with existing capacity constraints (LP1) do not take the relative link transmission durations into account, it severely overestimates the network capacity. Our new formulation (LP2) takes this factor into account and captures the fact that the capacity of the random-access network decreases as γ increases.

IX. CONCLUSIONS

We initiate the capacity analysis of asynchronous random-access wireless networks. Through the rigorous notion of competitive analysis, we precisely quantify the gap in network capacity due to the use of random-access scheduling as opposed to optimal scheduling, and design novel random-access strategies which achieve this limit. A key intuition which is an outcome of our work is that random-access scheduling is generally more competitive when packet sizes for each link is chosen in proportion to the link-capacity. Motivated by this intuition, we conjecture that random-access with proportional packet-sizes is α -competitive w.r.t. any other random-access strategy, for some constant $\alpha > 0$. Settling this conjecture is an interesting open problem. Another significant avenue for future research is the capacity analysis of 802.11 and related protocols which employ exponential back-off, RTS-CTS based collision avoidance schemes, and other sophisticated features.

ACKNOWLEDGMENTS

The authors wish to thank the anonymous reviewers for their insightful comments. D. Chafekar, V.S.A. Kumar, D. Levin, M.V. Marathe, and A. Srinivasan were supported in part by NSF Award CNS-0626964. D. Levin and A. Srinivasan were also supported in part by NSF ITR Award CNS-0426683. D. Levin was further supported by a UMD CSD Dean's fellowship. Part of this work was done while A. Srinivasan was on sabbatical at the Network Dynamics and Simulation Science Laboratory of the Virginia Bioinformatics Institute, Virginia Tech.

REFERENCES

- [1] M. Alicherry, R. Bhatia, and L. E. Li. Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 58–72, New York, NY, USA, 2005. ACM Press.
- [2] H. Balakrishnan, C. Barrett, A. Kumar, M. Marathe, and S. Thite. The Distance 2-Matching Problem and Its Relationship to the MAC Layer Capacity of Adhoc Wireless Networks. *special issue of IEEE J. Selected Areas in Communications*, 22(6):1069–1079, 2004.
- [3] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, 2000.
- [4] L. Bononi, M. Conti, and E. Gregori. Runtime optimization of IEEE 802.11 wireless lans performance. *IEEE Transactions on Parallel and Distributed Systems*, 15(1):66–80, 2004.
- [5] F. Cali, M. Conti, and E. Gregori. Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit. *IEEE/ACM Trans. Netw.*, 8(6):785–799, 2000.
- [6] L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Found. Trends Netw.*, 1(1):1–144, 2006.
- [7] M. Heusse, F. Rousseau, R. Guillier, and A. Duda. Idle sense: an optimal access method for high throughput and fairness in rate diverse wireless lans. *SIGCOMM Comput. Commun. Rev.*, 35(4):121–132, 2005.
- [8] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *Proceedings of the 9th annual international conference on Mobile computing and networking*, pages 66–80. ACM Press, 2003.
- [9] C. Joo and N. B. Shroff. Performance of random access scheduling schemes in multi-hop wireless networks. To appear in *Infocom*, 2007.
- [10] M. Kodialam and T. Nandagopal. Characterizing achievable rates in multi-hop wireless mesh networks with orthogonal channels. *IEEE/ACM Trans. Netw.*, 13(4):868–880, 2005.
- [11] M. Kodialam and T. Nandagopal. Characterizing the capacity region in multi-radio multi-channel wireless mesh networks. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 73–87, New York, NY, USA, 2005. ACM Press.
- [12] A. Kumar, E. Altman, D. Miorandi, and M. Goyal. New insights from a fixed point analysis of single cell IEEE 802.11 WLANs. In *INFOCOM*, pages 1550–1561, 2005.
- [13] V. S. A. Kumar, M. V. Marathe, S. Parthasarathy, and A. Srinivasan. Algorithmic aspects of capacity in wireless networks. In *ACM SIGMETRICS*, pages 133–144, 2005.
- [14] X. Lin and S. B. Rasool. Constant-time distributed scheduling policies for ad hoc wireless networks. In *Proceedings of IEEE Conference on Decision and Control*, 2006.
- [15] V. Ramaiyan, A. Kumar, and E. Altman. Fixed point analysis of single cell IEEE 802.11e WLANs: uniqueness, multistability and throughput differentiation. In *SIGMETRICS '05: Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 109–120, New York, NY, USA, 2005. ACM Press.
- [16] G. Sharma, A. Ganesh, and P. Key. Performance analysis of contention based Medium Access Control protocols. In *Proceedings of IEEE INFOCOM*, 2006.
- [17] G. Sharma, R. R. Mazumdar, and N. B. Shroff. On the complexity of scheduling in wireless networks. In *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, pages 227–238, New York, NY, USA, 2006. ACM Press.
- [18] W. Wang, X.-Y. Li, O. Frieder, Y. Wang, and W.-Z. Song. Efficient interference-aware TDMA link scheduling for static wireless networks. In *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, pages 262–273, New York, NY, USA, 2006. ACM Press.