Shaping Throughput Profiles in Multihop Wireless Networks: A Resource Biasing Approach

Sumit Singh, *Member, IEEE,* Upamanyu Madhow, *Fellow, IEEE,*
Elizabeth M. Belding, *Senior Member, IEEE,*

Abstract—A fundamental question in multihop wireless network protocol design is how to partition the network's transport capacity among contending flows. A classically "fair" allocation leads to poor throughput performance for all flows because connections that traverse a large number of hops (i.e., long connections) consume a disproportionate share of resources. However, naïvely biasing against longer connections can lead to poor network utilization, because a significantly high fraction of total connections are long in large networks with spatially uniform traffic. While proportional fair allocation provides a significant improvement, we show here that there is a much richer space of resource allocation strategies for introducing a controlled bias against resource-intensive long connections in order to significantly improve the performance of shorter connections. Specifically, mixing strongly biased allocations with fairer allocations leads to efficient network utilization as well as a superior tradeoff between flow throughput and fairness. We present an analytical model that offers insight into the impact of a particular resource allocation strategy on network performance, taking into account finite network size and spatial traffic patterns. We point to protocol design options to implement our resource allocation strategies by invoking the connection with the well-studied network utility maximization framework. Our simulation evaluation serves to verify the analytical design prescriptions.

Index Terms—Resource allocation, multihop wireless networks, proportional fairness, resource biasing, network utility maximization

1 Introduction

In this paper, we study resource sharing strategies for partitioning the transport capacity of static multihop wireless networks. The throughput performance for these networks is determined by how the transport capacity is divided up among competing flows. Consider a network of N identical nodes with fixed link rates and spatially uniform traffic, in which each node chooses a destination uniformly at random from among the other nodes. For such a network, Gupta and Kumar [1] showed that the achievable per-flow throughput diminishes approximately as $O(\frac{1}{\sqrt{N}})$. Intuitively, this negative scalability result follows from the fact that for a network deployment area A, the total transport capacity only increases as $O(\sqrt{AN})$, whereas the number of hops to the destination and hence the amount of network resources required per-connection scales up as $O(\sqrt{N})$. This per-flow throughput analysis implicitly assumes that the resource allocation to each connection is throughput-fair across the network. Connections that traverse a larger number of hops (longer connections) therefore consume significantly more network resources than those traversing fewer hops (shorter connections) to achieve the same end-to-end throughput.

- *S. Singh was and U. Madhow is with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106, USA (email: sumit@ece.ucsb.edu).*
- *E. Belding is with the Department of Computer Science, University of California, Santa Barbara, CA 93106, USA.*

A natural question that arises from the preceding observation is the following: what if we sidestep assumption of throughput-fair allocation by intelligently biasing network resource allocation against resource-intensive longer connections? Shorter connections could then achieve a significantly higher throughput at the cost of a controlled performance degradation for the longer connections. In this paper, we introduce a rich class of resource allocation strategies that embody this design philosophy. Our goal is to provide the network designer with a flexible means of trading off throughput performance seen by long and short connections in a manner consistent with network design goals.

Resource sharing strategies from prior literature that fall within our framework are proportional fairness [2], and the more general (p, β)-proportional fairness [3]. However, we show that it is possible to design the resource allocation strategy to choose from among a much larger class of flow throughput profiles (i.e., flow throughput versus the number of hops the flow traverses) while maintaining efficient network utilization. For example, it is possible to significantly improve the performance of shorter connections beyond that attained by proportional fairness, while minimally degrading the performance of longer connections. Alternatively, we can improve throughput for both short and long connections relative to proportional fairness, at the expense of slightly lower throughput for flows traversing a moderate number of hops.

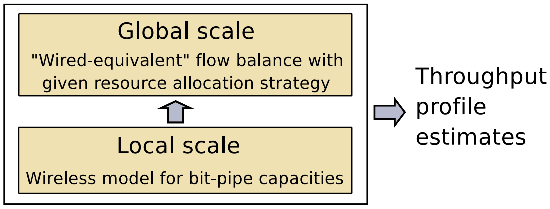

Fig. 1. Two-scale model for resource allocation over large multihop wireless networks.

While it is intuitively clear that taking away a small amount of resources from long connections can significantly improve performance for short connections, naïvely biasing against long connections can lead to poor performance. In a large network with spatially uniform traffic, a large fraction of total connections are long, so that there are not enough short connections to take advantage of the transport capacity that is released if we strongly bias against long connections. We draw upon this key observation to propose a new class of resource allocation strategies called *mixed-bias* strategies, in which a fraction of the total available network resource is allocated via an allocation strategy that is strongly biased against long connections, with the remainder allocated in an unbiased (e.g., max-min fair) or mildly biased (e.g., proportional fair) manner. Mixed-bias strategies enable the desired flexibility in shaping throughput profiles while maintaining efficient network utilization.

Our second contribution is an analytical model that provides quick estimates of the throughput profile achieved by a given resource allocation strategy. This model serves as a tool for exploring the rich design space of mixed-bias strategies in shaping throughput profiles that are consistent with network performance goals. The foundation for our analytical framework is a novel two-scale model for large multihop wireless networks (see Fig. 1). At the *global* scale, the decay of transmit power with distance and the broadcast nature of wireless medium (which limits spatial reuse), necessitate the use of short distance hops. Therefore, wireless plays a fundamental role only in influencing the network topology. At this scale, elementary flow calculations that treat node bandwidth shares (or data transfer capacity) as bit pipes suffice to analyze the impact of various resource allocation strategies. For instance, for throughput-fair scheduling at the nodes, Gupta and Kumar's negative scalability results follow from such flow analysis, as do results for different biased resource allocation strategies. The effective node data transfer capacities used in the preceding "wired-equivalent" global scale calculations are obtained from a detailed analysis of the wireless medium at the *local* scale, which takes into account medium access control, transceiver capabilities, and the radio propagation characteristics. While our approach is based on flow balance at a "typical" node in a large network, it is also refined to account for the effect of finite network size and spatial traffic patterns.

We verify the design prescriptions obtained from our analytical model using simulations of IEEE 802.11-based multihop networks. These simulations incorporate the effects of complex node interactions at different layers of the wireless network stack that are not explicitly modeled in our analysis. The performance trends from the simulations demonstrate that the two-scale model, despite its simplicity, captures the fundamental design tradeoffs.

Our third contribution is to show that the throughput profiles attained via application of mixed-bias resource allocation can be obtained as solutions to network utility maximization (NUM) [2] sub-problems in the following idealized setting: each bias corresponds to a different choice of concave utility function, a fixed fraction of network resources is allocated to each bias, and for each sub-problem, the shadow price of each bit-pipe in the network is the same (approximately satisfied for large, heavily loaded networks). This connection with NUM opens the way to leveraging the extensive literature on decentralized NUM-based protocols for implementing the class of resource biasing strategies proposed here. While detailed design and investigation of such protocols is beyond the scope of the present paper, we point out that such implementations would consider a single NUM for a given mixed-bias strategy, that there is no need to enforce fixed fractional resource allocations for each bias, and that in practice, the link shadow prices and resources allocated would vary with network topology and spatiotemporal traffic fluctuations. The utility functions for such NUMs extend the class of utility functions generally considered for resource allocation [3].

To summarize, our main contributions are as follows:

- We propose *mixed-bias* strategies that blend strongly biased resource allocation with fairer allocation strategies. We show that mixed-bias strategies achieve a superior flow throughput profile and efficient network utilization that cannot be attained by the individual strategies in the mixture.
- We present a two-scale analytical model that provides insight into the impact of a particular resource allocation strategy on network performance, in a manner that captures the effect of finite network size and spatial traffic patterns. We present packet-level simulations for an example application to verify the performance predictions from our analytical model.
- We show that mixed-bias resource allocation can be mapped to appropriately defined network utility maximization problems, which points the way for future work translating our mixed-bias approach to implementable network protocols.

The rest of the paper is organized as follows. Section 2 discusses the related work. Section 3 presents our resource allocation framework and describes mixed-bias allocation policies. In Section 4, we present an example application of our resource biasing framework and present a simulation evaluation. In Section 5, we establish the connection between our resource biasing strategies and a NUM framework, which shows that our resource allocation policies are amenable to distributed protocol implementations. We conclude with a summary of our findings in Section 6.

2 Related Work

We first introduced the resource biasing strategy and the two-scale model discussed in this paper in a conference publication [4]. In the present paper, we expand upon [4], providing more detailed analysis and simulation results, and making the connection to network utility maximization (NUM).

Most prior work on resource allocation in communication networks is based on a network utility maximization framework [2]: a user (i.e., flow) i attaining rate x_i obtains utility $U(x_i)$, where $U(\cdot)$ is a strictly concave, non-decreasing, continuously differentiable function, and the goal of resource allocation is to maximize the sum of utilities, subject to network resource constraints. Consider a network with a set L of resources or logical links with capacities $R = (r_l, l \in L)$, and a set of flows F. The basic NUM problem formulation for a wired network is as follows [2]:

$$\max \sum_{i \in F} U(x_i) \tag{1}$$

subject to

$$AX \le R \quad \text{over} \quad X \ge 0 \tag{2}$$

where $X = (x_i, i \in F)$ is the vector of user (flow) rates, $A = (A_{li}, l \in L, i \in F)$ is a $0-1$ routing matrix where $A_{li} = 1$ if flow i uses link l, $A_{li} = 0$ otherwise. The above formulation is a convex optimization problem, given the concavity of the objective function (1) and the compact feasible region (2). In addition to convexity, typical approaches to solving NUM problems identify and exploit a decomposition structure embedded in the problem to design decentralized algorithms that converge to the global optimum. The solution algorithms are cast into network protocols that can span multiple layers of the network stack.

A utility function can be interpreted as a knob to trade off the network throughput performance versus fairness among competing flows. Reference [3] proposes the following class of utility functions which have been commonly considered in the subsequent work on NUM.

$$U(x, \beta) = \begin{cases} (1-\beta)^{-1} x^{1-\beta} & \beta \ne 1 \\ \log(x) & \beta = 1, \end{cases} \tag{3}$$

where $x \ge 0$ is flow rate and parameter $\beta > 0$ captures the degree of fairness, where a higher β implies a more concave $U(x)$, and hence a fairer allocation. $U(x, \beta)$ maps to sum throughput maximization for $\beta \to 0$, proportional fairness [2] for $\beta = 1$, potential delay minimization [5] for $\beta = 2$, and approaches max-min fairness [6] for $\beta \to \infty$.

Following the seminal work of Kelly [2], there has been extensive research on NUM in the context of resource allocation problems in wireline networks, and more recently, wireless networks. Comprehensive surveys of work on the applications of the NUM framework can be found in [7]–[9]. For multihop wireless networks, [10], [11] propose congestion control algorithms based on NUM formulations. More recent papers [12]–[17] consider joint congestion control and scheduling problem to design cross-layer resource allocation algorithms using the NUM framework. Most of these papers basically employ combinations of ideas from distributed wireline congestion control (e.g., [2] and follow up work), and queue-lengths based wireless scheduling or random access MAC (many of which are inspired from [18] and its extensions).

Because of our interest in shaping throughput profiles, our initial approach in our conference publication [4] was different from the NUM framework, in that we explicitly assumed that the resource allocation for a flow is proportional to a specific decreasing function of the number of hops traversed by a flow, and then computed the proportionality constant via flow balance. By connecting these biasing strategies to the NUM framework in this paper, we show that the network designer can work with a much broader class of utility functions than that considered in the prior work on NUM. In turn, we expect the protocol design for implementing our resource biasing strategies will be able to draw upon the growing literature on NUM-based resource allocation over multihop wireless networks [10]–[17], [19].

Other relevant prior work includes a discussion of the tradeoff between network capacity and fairness for heterogeneous multihop wireless networks [20], which concludes with advocacy of proportional fairness as a reasonable compromise between rate efficiency and fairness, and observes that max-min fairness yields poor performance because all flows obtain essentially the same rate as the minimum rate flow. Reference [21] presents a realization of proportional fairness based on end-to-end rate control for a connection based on the number of hops it traverses. Reference [22] presents a simple average-based model to analyze the unbiased resource allocation problem for ad-hoc wireless networks and derives capacity scaling laws for random as well as spatially localized traffic patterns.

3 Resource Allocation Framework

We now introduce our model for the resource allocation framework. We first discuss the global scale analysis, which essentially involves flow balance in a wired-equivalent network whose capacities would be set by a local scale analysis.

For a uniform traffic model, each node generates a persistent flow, choosing a destination at random from among all other nodes. Our framework also accommodates spatially localized traffic distributions in which the probability of choosing a destination depends on the number of hops from the source. The biased resource allocation problem is to assign bandwidth to each flow as a function of its resource requirements (e.g., number of hops it traverses) such that the node data transfer capacity constraints are honored. Since it is intractable to identify bottleneck links for randomized traffic patterns, we apply an averaged analysis at the global scale to do flow balance for a "typical" node of the following form:

$$C \geq \sum_{h=1}^{h_{max}} n(h)c(h), \tag{4}$$

where C is the available average data transfer capacity or bandwidth share of a node (in bits per-second), $n(h)$ is the average number of h-hop connections traversing the node, $c(h)$ is the effective data rate assigned to h-hop flows, and h_{max} is the maximum number of hops. Note that $c(h) \leq s$, where s is the maximum stable data rate of a link. Both C and s are determined via a detailed local scale analysis of wireless medium access control and physical layer properties: s can be considered as the data transfer capacity of a node's wireless contention region, whereas C is a node's effective, average share of s. An alternative way to interpret equation (4) is to multiply both sides by N, the number of nodes in the network. Then the equation points to the problem of partitioning the aggregate one hop capacity of the network among competing flows. For an N-node network, $h_{max} = O(N)$ for a linear topology, and $h_{max} = O(\sqrt{N})$ for a two-dimensional grid.

For a uniform traffic model, a typical node is more likely to see a longer flow since such a flow traverses more hops / nodes, so that $n(h)$ increases with h (the next two sections show concrete examples for one and two dimensional networks). It therefore makes sense to make $c(h)$ a decreasing function of h, in order to ensure that the sum in (4) does not exceed C even as the network size (and hence h_{max}) becomes large. A throughput profile refers to the dependence of $c(h)$ with h. Our goal is to explore the design space of throughput profiles that are feasible, in terms of meeting the link balance condition (4).

For simplicity of exposition, we first illustrate our framework with some quick global scale computations for a one-dimensional network in Section 3.1. We then provide a more comprehensive analysis for a two-dimensional network which models the IEEE 802.11 network to be simulated.

3.1 Computations for a One-Dimensional Network

Consider a directed linear network with N nodes (arranged as a ring to avoid edge effects: see Fig. 2) and $C = 1$. For uniform traffic, the probability of a connection initiated by any given node traversing h hops is $p(h) = \frac{1}{N-1}$, $h = 1, ..., N-1$, since any of the other $N-1$ nodes in the network can be chosen with equal probability. On average, there are $Np(h)$ flows traversing h hops each in a network of N nodes (as each node originates a flow). Therefore, the expected total number of hops traversed by h hop flows in the network is given by $Np(h)h$. This implies that the average number of h-hop connections traversing a node is given by $n(h) = \frac{Np(h)h}{N} = hp(h)$. Thus, flow balance at a typical node becomes (assuming full utilization):

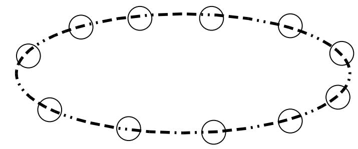

Fig. 2. Linear topology.

$$\sum_{h=1}^{N-1} hp(h)c(h) = 1 \tag{5}$$

For a throughput-fair allocation, the bandwidth allocated to an h-hop connection is independent of h, so that $c(h) = \lambda$. We show that the preceding throughput profile maps to max-min fairness in Section 5. Solving (5), we obtain $c(h) = \lambda \approx \frac{2}{N}$. The sum throughput $T_{sum} = N\sum_h c(h)p(h)$, summed over all connections, is approximately 2.

Now consider the case when the bandwidth assigned to a connection is inversely proportional to the number of hops it traverses; i.e., $c(h) = \frac{\lambda}{h}$. Here λ denotes the maximum allowed source application data rate (in bits per-second). We show that this allocation approximates proportional fairness in Section 5. Plugging this into (5), we obtain that $\lambda \approx 1$, so that $c(h) \approx \frac{1}{h}$. The sum throughput $T_{sum} \approx \log N$. Thus, we have significantly improved both the sum throughput and the performance of shorter connections, with only a factor of two loss in the throughput obtained by the longest connections. One can show using typical link flow analysis that the throughput of a h-hop connection is approximately $\frac{1}{h}$, with the sum throughput of $\log N$, and $\lambda = O(1)$. Comparing with the max-min fair allocation, we realize that biasing against the long connections via proportional fairness results in a win-win strategy: it leads to far better performance for the shorter connections, while reducing the throughput attained by the long connections only by roughly a factor of two. The cumulative effect of this is the improvement of the total throughput by a factor of $\frac{\log N}{2}$ compared to the max-min fair allocation.

The pitfalls of naïve biasing: Given the improved performance tradeoff offered by proportional fairness, it is natural to ask if we can get still larger improvements by biasing even more severely against the long connections. Consider a more severely biased allocation $c(h) = \frac{\lambda}{h^2}$. Substituting into (5), it can be shown that there is indeed a gain in sum throughput, but we require that $\lambda = O(\frac{N}{\log(N)})$. That is, the maximum flow throughput must scale up with network size, which is clearly not compatible with the finite capacity of one. Basically, there

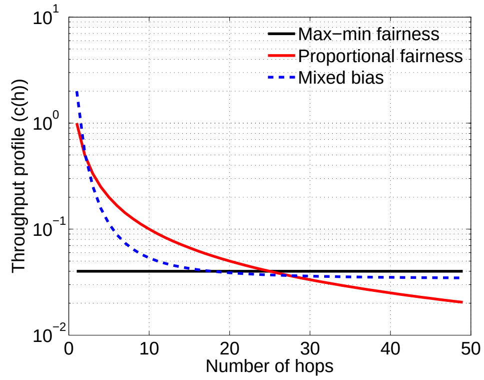

Fig. 3. Throughput profiles for mixed biasing over 50 node linear (ring) topology.

are not enough short flows to take advantage of the transport capacity released by such a strong bias. If we now impose the constraint that λ cannot scale up with N, we find (details are provided for a two-dimensional model later) that the network is poorly utilized.

The promise of mixed biasing: The problem with the preceding biasing strategy was that the short connections would have to send at too high a rate (more than the maximum stable link rate s) to take advantage of the bias. However, if we only allocate a small fraction of the available capacity to this biasing strategy, then we can eliminate this problem. For example, consider a *mixed-bias* strategy of the form $c(h) = \lambda_1 + \frac{\lambda_2}{h^2}$; this mixes max-min fairness with a bias stronger than proportional fairness. Let us assume that we use a proportion α_1 of the node data transfer capacity to implement max-min fairness, and a proportion $\alpha_2 = 1 - \alpha_1$ to implement the stronger bias such that flow balance decomposes into two separate equations:

$$\sum_{h=1}^{N-1} hp(h)\lambda_1 = \alpha_1, \text{ and } \sum_{h=1}^{N-1} hp(h)\frac{\lambda_2}{h^2} = \alpha_2.$$

Now, by choosing α_2 to scale down with N, we can prevent λ_2 from scaling up with N. Specifically, we can set $\alpha_2 = a\frac{\log N}{N}$, which will result in $\lambda_2 \approx a$, which no longer scales up with N. The resulting throughput profile $c(h)$ is a convex combination of the throughput profiles from the two different strategies:

$$c(h) \approx \frac{a}{h^2} + \left(1 - a\frac{\log N}{N}\right)\frac{2}{N}. \tag{6}$$

Here, we constrain $a \leq \frac{s-2/N}{1-2\log(N)/N^2}$ so that the maximum throughput does not exceed the stable link rate s. We note that (6) provides a very different throughput profile than either max-min fairness or proportional fairness. In particular, the throughput seen by long connections is almost as good as that of max-min fairness (and hence is a factor of two better than that obtained from proportional fairness), while the throughput seen by short connections does not scale down as a function of N (and hence is better than max-min fairness). Fig. 3 illustrates example throughput profiles for a 50-node linear network with link data rate $s = 2$ and node bandwidth share $C = 1$. In short, mixed-bias strategies open up a large design space not covered by existing resource allocation strategies. We shall explore this in more detail in our discussion of two-dimensional networks.

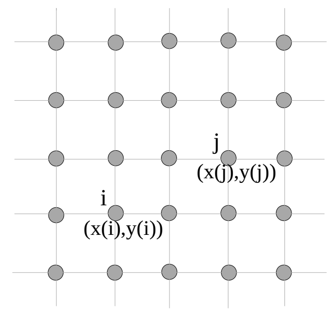

Fig. 4. Regular grid topology.

3.2 Two-Dimensional Networks

We now consider two-dimensional N node regular grid topology networks where the inter-node distance is such that direct communication is only possible between immediate neighbors on the grid. The hop distance $h_{(i,j)}$ between two nodes i and j with coordinates (x_i, y_i) and (x_j, y_j), respectively, is given by $h_{(i,j)} = |x_i - x_j| + |y_i - y_j|$, which equals the number of hops along the shortest path between the nodes (see Fig. 4).

We consider a power-law traffic model [23] where the probability that a node i will communicate with a node j, h hops away, is given by

$$p_{i,j} = \frac{q}{h^k}, \tag{7}$$

where $0 < q < 1$. Note that $p_{i,j} = p_{j,i} = p(h)$. Thus, the probability $p_H(h)$ that a pair of communicating nodes are h hops apart is given by

$$p_H(h) = \frac{m(h)}{h^k \sum_{r=1}^{h_{max}} \frac{m(r)}{r^k}}, \tag{8}$$

where $m(h)$ is the total possible number of h-hop flows in the network. For $k > 0$, this spatial traffic distribution models a spatially localized traffic pattern with degree of localization increasing with k. We first consider uniformly distributed traffic ($k = 0$) to study the effect of different resource biasing policies.

We focus on allocation of the available data transfer capacity C (in bits per-second) of a typical network node using link flow analysis to obtain the expected per-connection throughput. The underlying medium access control (MAC) layer (e.g., IEEE 802.11) implements max-min fair bandwidth allocation among contending nodes. The bias weight function $w(h)$ is defined as $w(h) = \frac{1}{h^b}$, where the bias exponent $b \geq 0$ determines the degree of bias. Let $n(h)$ denote the average number of h-hop flows passing through or initiated by a typical

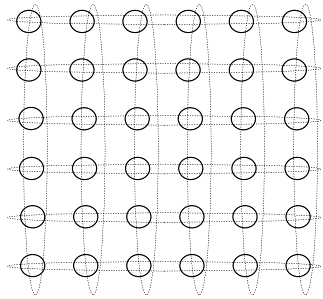

Fig. 5. Two-dimensional torus topology.

network node. We have

$$C \geq \sum_{h=1}^{h_{max}} n(h) \frac{\lambda}{h^b}. \quad (9)$$

Thus, the required λ for maximum utilization of the available capacity is given by $\lambda = C / \sum_{h=1}^{h_{max}} \frac{n(h)}{h^b}$.

We first consider a symmetric, regular grid (i.e., a 2-d torus, as shown in Fig. 5) topology to avoid edge effects, and focus on how the required λ scales with increasing network size as a function of the traffic model and the bias weight function. For a symmetric regular grid topology, $p_H(h) = 1/(h^{(k-1)} \sum_{r=1}^{h_{max}} \frac{1}{r^{(k-1)}})$. This is because the number of nodes located h hop away from any given node is $4h$, which yields the total possible number of h hop flows as $4Nh$. This is plugged in (8) to obtain $p_H(h)$. With the average number of h hop connections traversing a node given by $n(h) = hp_H(h)$, we have from (9),

$$\lambda = \frac{C}{\sum_{h=1}^{h_{max}} hp_H(h)/h^b} = \frac{C \sum_{h=1}^{h_{max}} 1/h^{(k-1)}}{\sum_{h=1}^{h_{max}} 1/h^{(k+b-2)}}. \quad (10)$$

For uniform traffic distribution ($k = 0$), we infer that with unbiased resource allocation ($b = 0$), $\lambda = O(1/\sqrt{N})$; for proportional fairness ($b = 1$), we have $\lambda = O(1)$; and for $b = 2$, $\lambda = O(\sqrt{N})$. Further calculations show that the required λ grows with the network size N for $b > 1$. However, since λ cannot exceed the maximum stable data rate s, this means that after the network size increases to a certain limit for $b > 1$, a node cannot fully utilize its fair share of bandwidth because the stronger bias prevents long flows from using the leftover capacity from the short flows. As in the linear network case, we infer that with proportional fair resource allocation, the per-flow throughput will be higher than the unbiased resource allocation for any N since the required λ is $O(1)$. Note that the preceding analysis also applies to infinite regular grid topologies where edge effects can be neglected.

Spatially localized traffic distribution: We now consider the effect of spatially localized traffic ($k > 0$). From (9) and (10), we find that for $k > 3$ in $p_H(h)$ defined in (8), the per-flow throughput stays constant, i.e., $O(1)$, rather than diminishing with increasing network size for unbiased resource allocation ($b = 0$). Even for traffic distributions with $2 < k \leq 3$, biased resource allocation with $b > 3 - k$ can ensure that per-flow throughput does not diminish. The observations about spatially localized traffic patterns and unbiased resource allocation are consistent with insights from [1], [22].

Finite, regular grid topologies: We now consider finite, regular grid topologies where the contention at the network edges is less than the core because there are fewer contending nodes at the edges; also, with a uniform traffic model where every node randomly picks a destination among other network nodes, the traffic volume increases towards the core of the network. Therefore, the network core acts as a bottleneck for network flows in most cases. Hence, we focus on the network-core to avoid offsets from the edge effects. We expect that the flow throughput results obtained from analysis of the network core are pessimistic estimates of the actual flow throughputs, especially for smaller networks (e.g., $N < 50$) where the edge effect dominates. Based on these observations, we modify (9) as follows:

$$C \geq \sum_{h=1}^{h_{max}} n_c(h) \frac{\lambda}{h^b}, \quad (11)$$

where $0 < \lambda \leq s$ and $n_c(h)$ denotes the average number of h-hop flows passing through or initiated by a typical network-core node. We evaluate $n_c(h)$ by averaging the values obtained from MatLab simulations of traffic instances over the network for 5000 seed values. Thus, we can calculate the maximum λ that satisfies (11) for a given bias exponent b. This λ value can be used to estimate the average per-flow throughput T of the network as a function of the network size, traffic distribution and the resource allocation policy as follows:

$$T = \sum_{h=1}^{h_{max}} p_H(h) \frac{\lambda}{h^b}. \quad (12)$$

Mixed-Bias Resource Allocation Policy: The underlying idea is to allocate a portion of the total available capacity at a node via a strongly biased policy, and allocate the rest employing a fairer policy. Assume that, of the total available node capacity C, we allocate αC via a resource allocation policy determined by weight function $w_1(h) = \frac{1}{h^{b_1}}$, and allocate the remaining bandwidth via weight function $w_2(h) = \frac{1}{h^{b_2}}$. We calculate the maximum λ_1 and λ_2 ($0 < \lambda_1, \lambda_2 \leq s$) that satisfy the following inequalities:

$$\alpha C \geq \sum_{h=1}^{h_{max}} n_c(h) \frac{\lambda_1}{h^{b_1}}, \quad (13)$$

$$(1-\alpha)C \geq \sum_{h=1}^{h_{max}} n_c(h) \frac{\lambda_2}{h^{b_2}}, \quad \lambda_1 + \lambda_2 \leq s \quad (14)$$

The average per-flow throughput in this case is given by:

$$T = \sum_{h=1}^{h_{max}} p_H(h) (\frac{\lambda_1}{h^{b_1}} + \frac{\lambda_2}{h^{b_2}}). \quad (15)$$

The parameters α, b_1 and b_2 allow us to span a large design space of resource allocation strategies, trading off throughput performance versus fairness. Note that an alternative formulation to (14) that ensures that any left-over capacity from the fraction of bandwidth allocated to strongly biased allocation in (13) is used up by the fairer allocation is given by $C \geq \sum_{h=1}^{h_{max}} n_c(h)(\frac{\lambda_1}{h^{b_1}} + \frac{\lambda_2}{h^{b_2}})$, where $\lambda_1 + \lambda_2 \leq s$. Using this approach, the choice of parameter α is less critical as α specifies the highest allowed fraction for the unfair strategy, and a high α value does not lead to resource under-utilization. In Section 5, we present a NUM formulation with a constraint based on the same idea. Since our intent is to first understand the effect of different parameters on the throughput profiles, we focus on (14) for our evaluations. We provide insight in to the effect of these parameters in the next section for an example application of our framework over IEEE 802.11 multihop networks.

4 Application to an IEEE 802.11 Multihop Network

We now apply our resource allocation framework to an IEEE 802.11 multihop network to illustrate how to determine global-scale parameters from local-scale analysis, which takes into account physical layer technology and medium access control. We then employ simulations with IEEE 802.11b parameters to determine how well the design prescriptions obtained by our analysis in Section 3 work.

For our local-scale analysis, Bianchi's saturation throughput analysis [24] is used to determine the data transfer capacity s of a contention region in an IEEE 802.11 network. Based on the observation that the IEEE 802.11 MAC protocol implements an approximation to max-min fair bandwidth allocation among contending neighbors [20], [25], a typical node's share of single hop capacity can be estimated as $C = s/n_c$, where n_c is the average number of contending nodes in a contention region at the network core. This value of C is now used in (11) in the global-scale analysis to estimate λ. This is then plugged into (12) to obtain the average per-flow throughput. Fig. 6 illustrates the predicted average per-flow throughput for an example mixed-bias allocation with $b_1 = 5$, $b_2 = 1$ (proportional fairness) and $\alpha = 0.2$, for the individual allocation policies in the mixture and also for max-min fair allocation. We observe that the mixed-bias allocation has a higher average flow throughput as compared to proportional fair and max-min fair allocations. Evaluation for different α values shows that increasing α to a higher value emphasizes strong bias whereas decreasing α to a smaller value such as 0.2 sways the throughput more towards the fairer allocation in the mixture.

We now employ simulations to obtain flow throughput profiles, and to evaluate the benefits of mixed-bias strategies relative to the individual strategies in the mixture, as well as to proportional fairness. We use the QualNet Network Simulator [26] over regular grid topology networks with IEEE 802.11b MAC/PHY, in basic access mode (RTS-CTS turned off). We present our results with static routing in order to isolate the effects of routing overhead on throughput. The results are averaged over 20 different seeds. Our application model comprises CBR flows of packet size of 1000 bytes, with data rates determined analytically for a given biasing strategy, as in Section 3.2 and enforced using source data rate control. Destinations are chosen at random based on the uniform traffic model.

While we fix the offered rate based on the analytical model in order to compare different resource allocation strategies, as we show in the next section, embedding mixed-bias resource allocation strategies within a utility function framework can provide the flexibility needed to arrive at decentralized adaptive implementations that allow nodes to dynamically tune their flow rates while implementing a biased resource allocation strategy. This can be used to alleviate congestion, to exploit spatial reuse opportunities, and more broadly, to react to specific network topologies and traffic patterns. In fact, embedding our biasing strategies on a NUM framework allows us to leverage protocol translations of jointly optimal cross-layer congestion control and scheduling algorithms, rather than restricting ourselves to flow rate control over a given MAC, as in this evaluation. Our aim in this section is to illustrate the promise of mixed biasing strategies even in this restrictive setting.

Fig. 7 illustrates the per-flow throughput performance for the example mixed-bias allocation considered in Section 3.2 ($b_1 = 5, b_2 = 1, \alpha = 0.2$) and compares it with that of the single bias policies in the mixture and also max-min fair allocation. The results are consistent with the analytical prediction that the mixed-bias allocation provides higher average flow throughput than max-min fair or proportional fair allocations. However, the average flow throughput values are lower, and decay faster with network size, than the analytical predictions shown in Fig. 6. This is because flows incur higher packet loss probabilities as they traverse a larger number of hops in a larger network, which results in an increase in the amount of wasted network resources. Another factor causing performance degradation that is not explicitly modeled in the analysis is the waste of medium access time due to the interaction between the IEEE 802.11 back-off mechanism and multihop packet forwarding [22].

Fig. 8 presents the scatterplots of individual flow throughputs as functions of the number of hops for a 144-node network over all simulation seeds. We consider a mixed-bias strategy, mixing a strong bias with proportional fairness. The strategy leads to significantly higher throughput for shorter connections than proportional fairness, while incurring virtually no degradation for the longer connections. Other desirable mixed-bias strategies that, for example, provide better performance for long connections than proportional fairness, can be obtained by mixing a strongly biased strategy with max-min

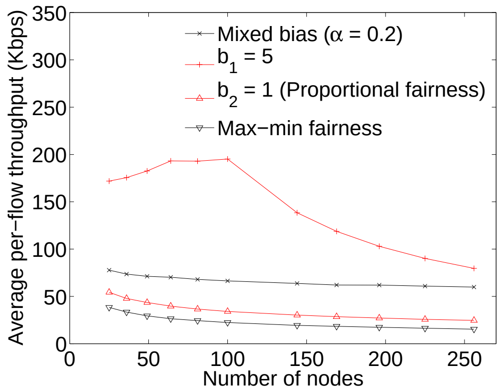

Fig. 6. Average per-flow throughput for mixed-bias with $b_1 = 5$, $b_2 = 1$, and $\alpha = 0.2$ from analysis.

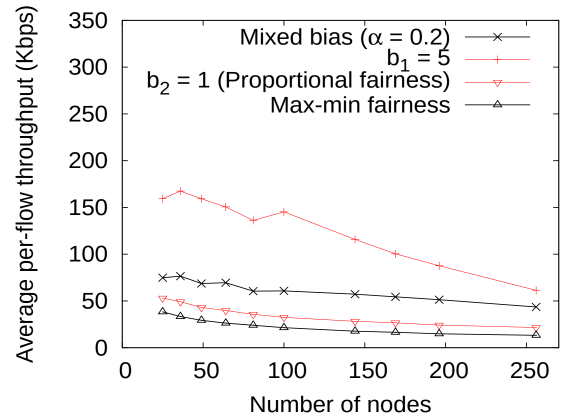

Fig. 7. Average per-flow throughput for mixed-bias with $b_1 = 5$, $b_2 = 1$, and $\alpha = 0.2$ from simulations.

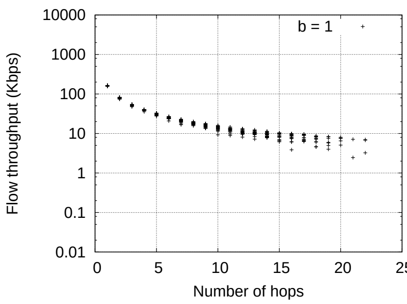

(a) Proportional fairness ($b = 1$).

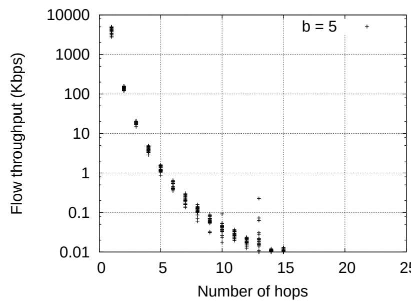

(b) $b = 5$.

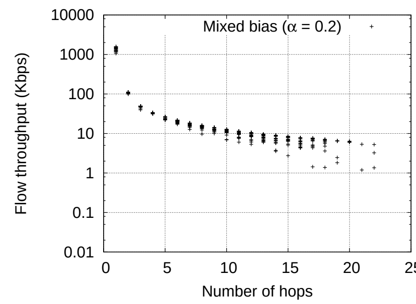

(c) Mixed ($b_1 = 5, b_2 = 1, \alpha = 0.2$).

Fig. 8. Scatterplots of flow throughputs for a 144-node network for mixed-bias allocations, proportional fairness and severe bias ($b = 5$).

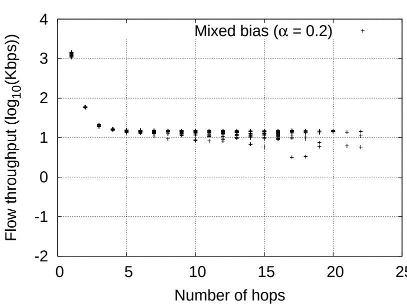

(a) Mixed ($\alpha = 0.2$).

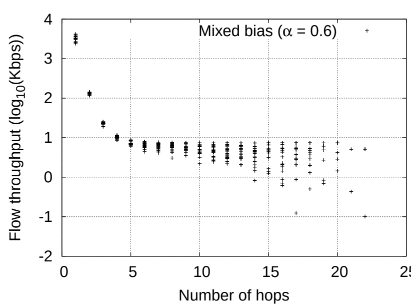

(b) Mixed ($\alpha = 0.6$).

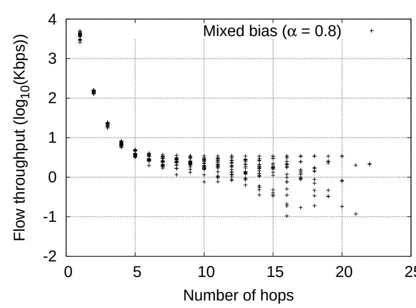

(c) Mixed ($\alpha = 0.8$).

Fig. 9. Scatterplots of flow throughputs for a 144-node network for mixed-bias allocation policies ($b_1 = 5, b_2 = 0$).

fairness (e.g., Fig. 9(a)). Note that the strongly biased strategy shuts out the long connections, and is therefore not a feasible choice on its own. The choice of the capacity fraction α and the strength b of the bias are other parameters at the disposal of the system designer to craft a desirable throughput profile.

We now investigate the effect of different values of the capacity fraction α allocated via high bias. Consider the mixed-bias allocation with ($b_1 = 5, b_2 = 0$). Fig. 10 illustrates the throughput performance obtained from the simulations for the mixed-bias policies for $\alpha = 0.2$, 0.6, and 0.8. The figure also plots single bias policies in the mixture and proportional fair allocation for reference. Figs. 9(b) and 9(c) present the scatterplots of individual flow throughputs as functions of the number of hops for a 144-node network over all simulation seeds, for the mixed-bias allocations with capacity fraction α = 0.6 and 0.8 (see Fig.9(a) for α = 0.2). A comparison between Figs. 9(b) and 9(c) illustrates the poor network utilization with a higher bias and high capacity fraction α. We see that allocation of 80% of the node capacity via a strong bias (i.e., $\alpha = 0.8$, $b_1 = 5$) results in decrease of

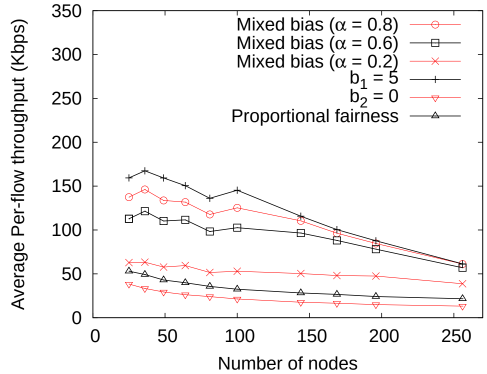

Fig. 10. Average per-flow throughput for mixed-bias with $b_1 = 5$, $b_2 = 0$, and $\alpha =$ 0.2,0.6, and 0.8 from simulations.

throughput achieved by long flows as compared to the case when $\alpha = 0.6$, but no apparent gain is observed in the throughput performance of the shorter flows. This demonstrates inefficient network utilization with a predominantly higher bias allocation for uniform traffic distribution. With different choices of capacity fraction α, it is possible to achieve a range of throughput profiles depending on the set of constituent biases in the mixture, the network size, and the traffic distribution.

5 Mapping to Network Utility Maximization

We now demonstrate that our resource biasing strategies are consistent with the network utility maximization (NUM) framework. Consider the convex optimization problem (1)–(2), for which a unique optimal solution exists [27]. From the Karush-Kuhn-Tucker conditions, the solution satisfies

$$U^{'}(x_i) = \sum_{l:l \in i} \mu_l \quad \forall i \in F \tag{16}$$

where $\mu = (\mu_l, l \in L)$ is a vector of Lagrange multipliers or shadow prices for the links, $\{l : l \in i\}$ denotes the set of links traversed by flow i, and X in (2) and μ satisfy

$$AX \leq R, \quad \mu \geq 0, \quad \text{and} \quad \mu^T(AX - R) = 0. \tag{17}$$

We now show that the resource allocation strategies discussed in Section 3 fall within the NUM framework in an idealized setting. We assume that all bit-pipe capacities in the global-scale wired equivalent model are approximately fully utilized: $AX \approx R$. We also assume that the aggregate shadow price for a flow is approximately proportional to the number of hops (bit-pipes) along its path, i.e., $\sum_{l:l\in i} \mu_l \approx h_i \mu_{avg}$, where $\mu_{avg} = \frac{\sum_l \mu_l}{|L|}$, $|L|$ is the cardinality of the set of bit-pipes. These assumptions provide insight into the performance in the core of large, heavily loaded networks, where we expect the traffic to be spatiotemporally uniform.

From (16), we have

$$U^{'}(x_i) = h_i \mu_{avg}, \quad \forall i \in F. \tag{18}$$

First consider the resource biasing throughput profile $x_i(h_i) = \lambda h_i^{-1}$, where $\lambda > 0$ is calculated via flow balance as described in Section 3. We now verify our claim that this throughput profile maps to a proportional fair allocation. We have $h_i = \frac{\lambda}{x_i}$, for which we find that

$$U(x_i) = k \log(x_i), \tag{19}$$

where $k = \lambda \mu_{avg}$ satisfies (18). This form of utility function $U(x_i)$ corresponds to proportional fairness. The extension for weighted proportional fairness (i.e., the utility function corresponding to each $i \in F$ is scaled with a weight $w_i > 0$) is straightforward.

Next, consider a biased allocation with $x_i(h_i) = \lambda h_i^{-b}$, where $\lambda, b > 0$. Hence, $h_i = \lambda x_i{}^{-\frac{1}{b}}$. We find that

$$U(x_i) = k \frac{x_i^{1-\beta}}{1-\beta}, \tag{20}$$

where $\beta = \frac{1}{b}$ and $k = \lambda \mu_c$, satisfies (18). Note that (20) is of the same form as the utility function for (p, β)-proportional fairness proposed in [3], with $\beta = \frac{1}{b}$.

Note that most prior work considering the class of utility functions proposed in [3] focus on $\beta \geq 1$, which maps to proportional fairness ($\beta = 1$), delay minimization ($\beta = 2$), and other utility functions with higher throughput fairness, approaching max-min fairness as $\beta \to \infty$. These functions map to our biased throughput profiles with bias exponent $1 > b > 0$, with $b \to 0$ tending to max-min fairness: $x_i(h_i) = \lim_{b\to 0} \lambda h_i^{-b} = \lambda$, where the rate λ per flow is derived via flow balance. as in Section 3. On the other hand, utility functions of the form (20) with $0 < \beta < 1$ map to strong bias against long connections ($b > 1$). Recall from Section 3 that allocating all network resources via such strongly biased policies can lead to poor network utilization for large networks, which is what motivated our proposed mixed-bias strategies. For our uniform traffic idealization, such mixed-bias allocations result from solving the following NUM problems:

$$\max \sum_{i \in F} U_1(x_{i,1}) + U_2(x_{i,2}) \tag{21}$$

subject to

$$AX_1 \leq \alpha R, \tag{22}$$

$$AX_2 \leq (1-\alpha) R, \tag{23}$$

where $U_1(.)$ is the utility function corresponding to the strongly biased allocation strategy (e.g., (20) with $0 < \beta < 1$), $U_2(.)$ corresponds to a fairer allocation such as max-min fairness or proportional fairness, $\alpha > 0$ represents the fraction of total resources allocated via the strongly biased allocation, and the flow rate vector $X = X_1 + X_2$ (where X_1, X_2 are constrained to be nonnegative). For example, mixing proportional fairness with a higher bias strategy with bias exponent $b = 5$ would lead to $U(x_i) = (\frac{(x_{i,1})^{0.8}}{0.8} + \log(x_{i,2}))$. The total rate for flow i is $x_i = x_{i,1} + x_{i,2}$. The form of the mixed biasing optimization problem is such that it can be decomposed into two parallel NUM problems in a straightforward manner, one with objective function $\sum_{i \in F} U_1(x_{i,1})$, under constraints (22), and the second with objective function $\sum_{i \in F} U_2(x_{i,2})$ and constraints (23). Thus, we have shown that the mixed-bias strategy corresponds to separate NUMs with a fixed resource partitioning governed by α.

Of course, the problem with the preceding mapping is that making the "right" choice for α is critical: it must be chosen such that, at least in our idealized model, full utilization is assured in (22) across biases, the aggregate shadow price of a flow is proportional to the number of hops. In practice, the network traffic would not be uniform, there would be edge effects, and it would be difficult to determine a "right" value of α. Thus, a more practical formulation is to couple the constraints in the NUM problems for different biases, so as to ensure full resource utilization even if the nominal value of α (which could potentially be based on analysis of our idealized model) is too large. Specifically, consider the following NUM:

$$\max \sum_{i \in F} U_1(x_{i,1}) + U_2(x_{i,2}) \tag{24}$$

subject to

$$AX_1 \leq \alpha R, \tag{25}$$

$$A(X_1 + X_2) \leq R. \tag{26}$$

We assume that the second utility function corresponds to a strategy such as proportional or max-min fairness which can fill up all network resources if needed (e.g., even if we set $\alpha = 0$). Thus, even if the value of α is too large and the strongly biased allocation cannot utilize all the network resources, the constraint (26) coupling the two biases implies that all network resources will be fully utilized. Note that, for a given value of α, the coupled NUM (24)-(26) has a larger feasible set, and hence at least as big a utility, as the decoupled NUMs (21)-(23).

Now that we have shown a mapping between the throughput profiles from our resource biasing strategies and the NUM framework, the extensive literature on decentralized, cross-layer algorithms for the NUM problems (see Section 2) can be leveraged to devise protocols that implement our resource biasing strategies. In practice, a decentralized NUM implementation would allow flexible resource partitioning (with an upper bound on the proportion of resources allocated via the highly biased strategy) and also different shadow prices to reflect the actual spatial distribution of traffic. Detailed investigation of design and evaluation of such NUM-based protocols is beyond the scope of this paper, but is an important topic for future work.

6 Conclusions

The resource biasing framework presented in this paper opens up a rich design space for sharing transport capacity in multihop wireless networks that goes beyond existing paradigms such as max-min fairness and proportional fairness by blending strongly biased and fairer allocations to get superior throughput profiles with high network utilization. Our two-scale model for multihop wireless networks yields quick performance estimates for different biasing strategies, providing a tool that allows system designers to tune parameters so as to shape throughput profiles while maintaining network efficiency. For 802.11 networks with fixed link speeds, the global scale, wired-equivalent model provides predictions of throughput profiles that match trends obtained by simulation. We have shown that, for an idealized setting of a network with uniform link speeds and traffic, mixed-bias resource allocations are solutions to parallel network utility maximization (NUM) problems, each maximizing a utility function corresponding to a particular bias, with a fixed fraction of network resources allocated to each bias. However, we note that practical implementations should be based on a coupled NUM in which a fair allocation capable of fully utilizing network resources is mixed with a highly biased strategy, setting an upper bound on the fraction of network resources that can be used by the latter. Decentralized implementations based on such a NUM can react to spatiotemporal variations in traffic and network topology. An important topic for future research, therefore, is to leverage the extensive research on NUM over wireless multihop networks (e.g., see [28], [12] for frameworks for joint congestion control and scheduling) to devise robust, decentralized protocols for implementing mixed-bias strategies.

Acknowledgments

This work was supported in part by the National Science Foundation under grants CNS-0347886 (CAREER award), CNS-0435527 and CNS-0520335 (under the Nets program), and by the Office of Naval Research under grant N00014-06-1-0066.

References

[1] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inform. Theory*, vol. 46, no. 2, pp. 388–404, 2000.

[2] F. Kelly, A. Maulloo, and D. Tan, "Rate control In communication networks: shadow prices, proportional fairness and stability," *J. Operational Research Soc.*, vol. 49, pp. 237–252, 1998.

[3] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 556–567, 2000.

[4] S. Singh, U. Madhow, and E. M. Belding, "Beyond proportional fairness: A resource biasing framework for shaping throughput profiles in multihop wireless networks," in *IEEE INFOCOM 2008 Mini Conference*, Phoenix, AZ, April 2008.

[5] L. Massouli and J. Roberts, "Bandwidth sharing: objectives and algorithms," *IEEE/ACM Trans. Netw.*, vol. 10, no. 3, pp. 320–328, 2002.

[6] L. Tassiulas and S. Sarkar, "Max-min fair scheduling in wireless networks," in *Proc. IEEE INFOCOM 2002*, New York City, NY, June 2002, pp. 763–772.

[7] R. Srikant, *The mathematics of Internet congestion control*. Birkhauser, 2004.

[8] X. Lin, N. B. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 8, pp. 1452–1463, 2006.

[9] M. Chiang, S. Low, A. Calderbank, and J. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 255–312, 2007.

[10] Y. Yi and S. Shakkottai, "Hop-by-hop congestion control over a wireless multi-hop network," *IEEE/ACM Trans. Netw.*, vol. 15, no. 1, pp. 133–144, 2007.
[11] M. Chiang, "Balancing transport and physical layers in wireless multihop networks: Jointly optimal congestion control and power control," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, pp. 104–116, 2005.
[12] U. Akyol, M. Andrews, P. Gupta, J. D. Hobby, I. Saniee, and A. L. Stolyar, "Joint scheduling and congestion control in mobile ad-hoc networks," in *Proc. IEEE INFOCOM 2008*, Phoenix, AZ, April 2008, pp. 619–627.
[13] A. Eryilmaz and R. Srikant, "Joint congestion control, routing, and MAC for stability and fairness in wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1514–1524, 2006.
[14] L. Bui, R. Srikant, and A. Stolyar, "Novel architectures and algorithms for delay reduction in back-pressure scheduling and routing," in *Proc. IEEE INFOCOM 2009 Mini Conference*, Rio de Janeiro, Brazil, 2009, pp. 2936 – 2940.
[15] M. Neely, E. Modiano, and C.-P. Li, "Fairness and optimal stochastic control for heterogeneous networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 2, pp. 396–409, April 2008.
[16] J. Lee, M. Chiang, and R. A. Calderbank, "Jointly optimal congestion and contention control based on network utility maximization," *IEEE Commun. Lett.*, vol. 10, no. 3, pp. 216–218, 2006.
[17] X. Lin and N. 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, 2006.
[18] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, 1992.
[19] L. Chen, S. H. Low, and J. C. Doyle, "Joint congestion control and media access control design for ad hoc wireless networks," in *Proc. IEEE INFOCOM 2005.*, vol. 3, Miami, FL, March 2005, pp. 2212–2222.
[20] B. Radunovic and J.-Y. L. Boudec, "Rate performance objectives of multihop wireless networks," *IEEE Trans. Mob. Comput.*, vol. 3, no. 4, pp. 334–349, 2004.
[21] A. Velayutham, K. Sundaresan, and R. Sivakumar, "Non-pipelined Relay Improves Throughput Performance of Wireless Ad-Hoc Networks." in *Proc. IEEE INFOCOM 2005*, Miami, FL, March 2005, pp. 477–490.
[22] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *Proc. Mobicom'01*, Rome, Italy, July 2001, pp. 61–69.
[23] N. Zhou, H. Wu, and A. A. Abouzeid, "The impact of traffic patterns on the overhead of reactive routing protocols," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 3, pp. 547–560, Mar. 2005.
[24] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535–547, March 2000.
[25] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. IEEE INFOCOM 2003*, vol. 2, San Francisco, CA, March-April 2003, pp. 836–843.
[26] (2005) QualNet network simulator, version 3.9. [Online]. Available: http://www.scalable-networks.com
[27] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
[28] A. Stolyar, "Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm," *Queueing Systems*, vol. 50, no. 4, pp. 401–457, 2005.

Sumit Singh received the bachelors degree in Electrical Engineering from the Indian Institute of Technology, Bombay, in 2002, and masters and Ph.D. degrees in Electrical and Computer Engineering from the University of California, Santa Barbara, in 2006 and 2009, respectively. He worked as a software engineer in the Network Systems Division at Samsung Electronics, Bangalore from 2002 to 2004. Sumit was a postdoctoral researcher with Prof. Upamanyu Madhow at UC Santa Barbara from 2009 to 2010, and is currently employed as a Wireless Architect at Moseley Wireless Solutions Group in Santa Barbara. His research interests lie in wireless networking, particularly medium access control, resource allocation, and quality of service in wireless networks.

Upamanyu Madhow is a Professor of Electrical and Computer Engineering at the University of California, Santa Barbara. His research interests are in communication systems and networking, with current emphasis on wireless communication, sensor networks and multimedia security.

He received his bachelor's degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1985, and his Ph. D. degree in electrical engineering from the University of Illinois, Urbana-Champaign in 1990. He has worked as a research scientist at Bell Communications Research, Morristown, NJ, and as a faculty at the University of Illinois, Urbana-Champaign.

Dr. Madhow is a recipient of the NSF CAREER award. He has served as Associate Editor for the IEEE Transactions on Communications, the IEEE Transactions on Information Theory, and the IEEE Transactions on Information Forensics and Security. He is the author of the textbook Fundamentals of Digital Communication, published by Cambridge University Press in 2008.

Elizabeth M. Belding is a Professor in the Department of Computer Science at the University of California, Santa Barbara. Elizabeth's research focuses on wireless networks, specifically multimedia, advanced service support, monitoring, and solutions for networking in under-developed regions. She is the founder of the Mobility Management and Networking (MOMENT) Laboratory (http://moment.cs.ucsb.edu) at UCSB. Elizabeth is the author of over 80 papers related to mobile networking and has served on over 50 program committees for networking conferences. Elizabeth served as the TPC Co-Chair of ACM MobiCom 2005 and IEEE SECON 2005, and the TPC Co-Chair of ACM MobiHoc 2007. She also served on the editorial board for the IEEE Transactions on Mobile Computing. Elizabeth is the recipient of an NSF CAREER award, and a 2002 Technology Review 100 award, awarded to the world's top young investigators. See http://www.cs.ucsb.edu/ ebelding for further details.