

A Tutorial on Cross-Layer Optimization in Wireless Networks

Xiaojun Lin, *Member, IEEE*, Ness B. Shroff, *Senior Member, IEEE* and R. Srikant, *Fellow, IEEE*

Abstract—This tutorial paper overviews recent developments in optimization based approaches for resource allocation problems in wireless systems. We begin by over-viewing important results in the area of opportunistic (channel-aware) scheduling for cellular (single-hop) networks, where easily implementable myopic policies are shown to optimize system performance. We then describe key lessons learned and the main obstacles in extending the work to general resource allocation problems for multi-hop wireless networks. Towards this end, we show that a clean-slate optimization based approach to the multi-hop resource allocation problem naturally results in a “loosely coupled” cross-layer solution. That is, the algorithms obtained map to different layers (transport, network, and MAC/PHY) of the protocol stack are coupled through a limited amount of information being passed back and forth. It turns out that the optimal scheduling component at the MAC layer is very complex and thus needs simpler (potentially imperfect) distributed solutions. We demonstrate how to use imperfect scheduling in the cross-layer framework and describe recently developed distributed algorithms along these lines. We conclude by describing a set of open research problems.

I. INTRODUCTION

Optimization based approaches have been extensively used over the past several years to study resource allocation problems in communication networks. For example, Internet congestion control can be viewed as distributed primal or dual solutions to a convex optimization problem that maximizes the aggregate system performance (or utility). Such approaches have resulted in a deep understanding of the ubiquitous Transmission Control Protocol (TCP) and resulted in improved solutions for congestion control [1]–[6].

The key question is whether such approaches can be applied to emerging multi-hop wireless networks to enable a clean-slate design of the protocol stack¹. Indeed there are unique challenges in the wireless context that do not allow a direct application of such techniques from the Internet setting. In particular, the wireless medium is an inherently multi-access medium where the transmissions of users interfere with each other and where the channel capacity is time-varying (due to user mobility, multipath, and shadowing). This causes

Xiaojun Lin and Ness B. Shroff are with the Center for Wireless Systems and Applications (CWSA) and the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA (email: {linx,shroff}@ecn.purdue.edu).

R. Srikant is with the Department of Electrical and Computer Engineering and Coordinated Science Laboratory, University of Illinois, Urbana, IL 61801-2307, USA (e-mail: rsrikant@uiuc.edu).

¹The notion of a clean-slate design becomes especially attractive for multi-hop wireless networks, where the burdens of legacy systems are far less than for the Internet.

interdependencies across users and network layers that are simply not present in their wireline counterparts. In spite of these difficulties, there have been significant recent advances that demonstrate that wireless resources across multiple layers (such as time, frequency, power, link data rates and end-user data rates), can be incorporated into a unified optimization framework. Interestingly, as will be described in detail in Section III, the solution of such an optimization framework will itself exhibit a *layered structure* with only a limited degree of cross-layer coupling.

We will illustrate the use of such an optimization approach for two classes of cross-layer problems, namely, the opportunistic scheduling problem in cellular (or access-point based single-hop networks), and the joint congestion-control and scheduling problem in multi-hop wireless networks. We will see that *convex programming* is an important tool for this optimization approach; in particular, Lagrange duality is a key tool in decomposing the otherwise complex optimization problem into easily-solvable components. However, we will also see that convex programming is often not enough. In fact, unlike their wireline counterparts, the essential features of many wireless cross-layer control problems are not convex. For example, due to interference, wireless networks typically require sophisticated “scheduling” mechanisms to carefully select only a subset of links to be activated at each time. In wireless networks, the capacity of each link depends on the signal and interference levels, and thus depends on the power and transmission schedule at other links. This relationship between the link capacity, power assignment, and the transmission schedule is typically non-convex. Therefore, the scheduling component needs to solve a difficult non-convex problem, and usually becomes the bottleneck of the entire solution.

These inherent non-convex features require that advanced techniques in addition to convex programming be used to satisfactorily solve the cross-layer control problem in wireless networks. In this tutorial, we will see a few examples where tools from convex programming, combinatorial optimization, stochastic stability, graph theory, large deviations, and heavy-traffic limits are used to obtain realistic and efficient solutions to the cross-layer control problem.

We acknowledge that cross-layer optimization has become a very active research area in the last few years. A comprehensive survey would be difficult due to the space constraints in this special issue. Hence, this tutorial is by no means an exhaustive survey of all subjects in cross-layer optimization.

Rather, our focus is to provide the readers with a sketch of the main issues, challenges, and techniques in this area, and also identify the main open problems to the community. For another survey in this area see [7].

The rest of the tutorial is organized as follows. Section II will begin with an exposition of the important problem of scheduling in cellular networks. Here, the emphasis is on incorporating physical layer channel information into the scheduling decision. In Section III, we investigate the joint congestion-control and scheduling problem in multihop wireless networks. The general formulation provided in Section III elegantly decomposes the cross-layer problem into a congestion control component and a scheduling component. However, due to non-convexity, the perfect scheduling component is usually very complex and difficult to implement in real networks. One approach to address this complexity is to use simpler (and potentially distributed) *imperfect* scheduling components in the cross-layer solution. However, the impact of these imperfect scheduling policies on the overall solution must be carefully studied. This is the subject of Section IV. In Section V, we will describe recent developments in obtaining imperfect distributed scheduling policies with provably achievable throughput bounds. We then conclude with a set of open problems.

II. OPPORTUNISTIC SCHEDULING FOR CELLULAR WIRELESS NETWORKS

In this section, we focus on the opportunistic scheduling problem in cellular networks (the results also apply to access-point based single-hop wireless networks). Over the last few years, this multi-user scheduling problem has received significant attention in both academia and industry. These scheduling schemes have been motivated by the unique features in wireless networks: scarce resources, mobile users, interference from other users in the network, and time-varying channel conditions (due to fading and mobility). Hence, good scheduling schemes in wireless networks should *opportunistically* seek to exploit channel conditions to achieve higher network performance. For example, consider a cellular network that consists of a base station and N users. Further, assume a time-slotted system and *downlink communications*, i.e., from the base-station to the users (receivers). Then the base-station can determine which user(s) to transmit to, based on the channel conditions. The idea is that transmissions to receivers with favorable channel conditions (e.g., with higher SINR) allows the base-station to transmit at a higher rate (using adaptive modulation and coding schemes) for a given target bit-error-rate. Thus, the base-station can *opportunistically* exploit the channel conditions to achieve higher network performance. It should be noted here that the idea of exploiting multi-user diversity is in contrast to traditional methods (e.g., spread spectrum, repetitive coding, power averaging, etc.), where the goal is to *smooth out* channel variability rather than to *exploit* it. Opportunistic scheduling achieves multi-user diversity gains because when users experiencing good channels are selected,

it enables the system to potentially operate close to its peak rather than average performance.

In [8], under an AWGN (additive white gaussian noise) model, it has been shown that the *sum capacity*² of a wireless system is maximized when only one user is selected to transmit at any given time. This result can be shown for either uplink or downlink assuming complete channel information at both the receiver and the transmitter. The user with the best channel condition is chosen for transmission. However, in a networking context, the difficulty with such a solution is that while it maximizes the overall throughput, it could result in significant unfairness among the users. For example, under such a scheme users that are close to the base-station may always be favored over those that are further away, resulting in potentially poor performance for certain users in the system. Such an approach is especially troubling for high-data-rate wireless users that may have stringent quality of service (QoS) requirements.

In order to address the above concerns, there have been several approaches to ensure fairness/QoS in a wireless context. For simplicity of presentation, we will focus on the *downlink*, i.e., base-station to user communication. We will overview opportunistic scheduling solutions that have been derived for both infinite- and finite-backlogged cases.

A. Infinite-Backlog Case

The infinite-backlog case is often studied in communication systems to evaluate access protocols and study their maximum achievable performance. It is also simple and results in a tractable solution that provides important insights. The objective in our context is to find a feasible scheduling policy Q that maximizes the overall system performance for given fairness/QoS requirements. A policy Q maps a vector $\vec{U} = [U_1(x_1), \dots, U_N(x_N)]$ to $Q(\vec{U})$, which is the index of the user selected for transmission. Here, x_i is the data rate transmitted to user i if it is selected for transmission, U_i is the utility function of user i , and $U_i(x_i)$ measures the value or benefit to user i of the receiving data rate x_i . Note that x_i is a function of the channel condition and the coding and modulation scheme used. There have been many scheduling schemes that address this problem [9]–[13]. Interestingly, most of these approaches result in an optimal solution that can be expressed in the form of simple myopic index policies given by:

$$Q^* = \underset{i=1, \dots, N}{\operatorname{argmax}} (\alpha_i U_i(x_i) + \beta_i), \quad (1)$$

where α_i and β_i are constants and can be viewed as Lagrange multipliers. For example, consider the following problem studied in [12], [13]:

$$\begin{aligned} \max_{Q \in \Theta} \quad & \sum_{i=1}^N E(U_i(x_i)) \\ \text{subject to} \quad & P\{Q(\vec{U}) = i\} \geq r_i, \quad i = 1, 2, \dots, N, \end{aligned} \quad (2)$$

where Θ is the set of all stationary scheduling policies, and r_i is the minimum fraction of time-slots assigned to user i

²The maximum total throughput that can be achieved in the system.

(i.e., fairness in time). Clearly, this is a non-linear optimization problem. However, it can be readily shown that policy (1) with $\alpha_i = 0$ and some appropriate choice of β_i is optimal (i.e., maximizes the expected system utility). The optimal policy (1) would be executed at each time-slot and β_i can be obtained through a stochastic approximation algorithm, as shown in [13]. Similarly, instead of temporal fairness, let us now consider other forms of QoS:

- *Fairness in utility*: Each user receives at least a fraction r_i of the aggregate utility value [11], [13].
- *Minimum data rate requirement*: Each user receives a minimum data rate of r_i bits per second [13].
- *Proportional Fairness*: Here, the objective is to achieve a solution that is proportionally fair, i.e., increasing the mean throughput of one user from the optimal level by $x\%$ results in a cumulative percentage decrease by greater than $x\%$ of the mean throughput of other users [14]. It turns out that such a solution is achieved when the optimization problem is to maximize the sum of the logarithms of the expected rates (or the product of the expected rates), i.e.,

$$\max_{Q \in \Theta} \sum_{i=1}^N (\log(E(x_i))).$$

In each of these cases it can be shown that the optimal solution will correspond to (1), now with $\beta_i = 0$. Thus, these results tell us that simple myopic index scheduling policies can be used to opportunistically improve the system performance in wireless networks.

B. Stability of Opportunistic Scheduling Schemes

The problem that we have described so far assumed an infinitely backlogged system with the objective of maximizing the aggregate system utility under QoS/fairness constraints. Another important class of problems deals with the development of opportunistic scheduling schemes with the intention of accommodating the maximum possible offered load on the system without violating stability or other QoS constraints. Here, the problem moves from the maximization of utility to stochastic stability. The work has largely been motivated by the seminal work on *throughput-optimal*³ scheduling [15]. This work shows that scheduling schemes that maximize the queue-weighted sum of the rates are throughput optimal. For the case of cellular networks, the scheduling scheme is of the following form:

$$Q^* = \operatorname{argmax}_{i=1, \dots, N} q_i x_i, \quad (3)$$

where q_i is the queue length of user i , and x_i is again the data rate transmitted to user i . While this scheduling solution does not account for fairness, it provides the important insight that queue-length information is critical in developing throughput optimal scheduling schemes. This idea has been further developed into a general class of queue-length based (or,

equivalently, delay-based) opportunistic scheduling schemes that focus on stability and throughput optimality [16]–[20].

For example, in [16], [17], simple index scheduling policies of the following form are shown to be throughput optimal:

$$Q^* = \operatorname{argmax}_{k=1, \dots, N} \alpha_k d_k x_k, \quad (4)$$

where α_k is a constant, d_k is the head-of-the-line packet delay at queue k , and as before x_k is the data rate of user k . In [17], [18], a related delay-based index policy that provides exponential weight to the delay (the so-called *exponential rule*) is shown to be throughput optimal.

Throughput optimal scheduling schemes have also been derived in [21], where the authors also incorporate flow-level dynamics into their model. In particular, the authors model users arriving to the system with a random amount of workload (e.g., a file size) and depart when this workload has been transmitted. Recently, in [22], [23] the authors have attempted to characterize the impact of different forms of scheduling on stability and QoS using techniques from large-deviation and heavy traffic limits. The key results from these works emphasize the importance of queue-length based (QLB) scheduling (e.g., in the form of (3) and (4)) for finite-backlogged systems when there are delay constraints. Under certain conditions, it can be shown that, when there are delay constraints, the network throughput of QLB policies is larger than policies for which queue-length information is not taken into account, e.g., (1). Moreover, for a given delay violation constraint, when the number of users in the system increases, the total network throughput under policy (1) initially increases and then eventually decreases to zero, but not so under the QLB policy. This should not be entirely surprising since index policies of the form (1) are agnostic to the delays incurred for different users and may not serve users whose queues are building up fast enough to remain within a delay violation probability.

C. Limitations and Lessons Learned

Thus far, we have focused on opportunistic scheduling solutions for cellular systems. For such systems, one can often find simple myopic index policies that are optimal and easy to implement. However, scheduling cannot address the problem of ensuring that the system is operating in a stable or feasible regime. Hence, while opportunistic scheduling expands the capacity region over its non-opportunistic counterparts, it may be difficult to utilize this gain if we are unable to operate the system close to the boundary of the capacity region. For example, if one were to make a conservative estimate of the boundary of the capacity region, and traffic were injected into the system based on this conservative estimate, the opportunistic gains may never be realized. Thus, it is imperative that one solves the problem of determining the rates to be injected into the network (i.e., congestion control to be discussed in Section III) jointly with which user(s) to be scheduled for transmission (i.e., opportunistic scheduling).

Further, when we move from cellular to multi-hop wireless systems, we encounter other difficulties that need to be

³A scheduling scheme is said to be throughput optimal if it stabilizes the system whenever any other feasible scheduler can stabilize the system.

addressed. For example, how should one determine the end-to-end data rates for the users? When should a given link be activated in the network? What should be the forwarding rate of each link along with its power allocation, coding and modulation schemes? How can one ensure that the rate provided by the links is enough to support the end-to-end rate of all users? The potential state space in a moderately sized multi-hop network could be quite large, so can one develop low-complexity solutions to these problems? Perhaps, most importantly, how does one go about developing efficient distributed solutions to these problems?

In the next few sections, we will describe some of the recent research developments that have taken place in addressing the above problems.

III. CROSS-LAYER CONGESTION CONTROL AND SCHEDULING FOR MULTI-HOP WIRELESS NETWORKS

In this section, we study the following problem in multi-hop wireless networks: how does one jointly choose the end-to-end data rate of each user and choose the schedule for each link? (Here, as in Section II, we will use the term *schedule* to refer to the joint allocation of resources at MAC/PHY layers, which include modulation, coding, power assignment and link schedules, etc.) As we will see, the solution to this problem is obtained by choosing an appropriate congestion control algorithm to regulate the user data rates and a scheduling policy which is a modification of the queue-length based scheduling algorithm in the previous section. While congestion control has been studied extensively for *wireline* networks [1]–[6], these results cannot be applied directly to multihop wireless networks because the link capacity in multihop wireless networks varies and depends on the scheduling policies at the underlying layers. There have been attempts to solve this cross-layer control problem using a “layered” approach [24]–[27]. The approach is to find a feasible *rate region* that has a simpler set of constraints similar to that of wireline networks and then develop congestion controllers that compute the rate allocation within this simpler rate region. Unfortunately, for general network settings, it is not always possible to find such a simpler rate region. Further, because the rate region reduces the set of feasible rates that congestion control can utilize, the layered approach results in a *conservative* rate allocation.

The general cross-layer solutions for jointly optimizing congestion control and scheduling have recently been developed by a number of researchers [20], [28]–[34]. In this section, we will review two types of formulations and solutions that can potentially be used for online implementation.

A. The Model

We consider a multi-hop wireless network with N nodes. Let \mathcal{L} denote the set of links (i.e., node pairs) (i, j) such that the transmission from node i to node j is allowed. Due to the shared nature of the wireless media, the data rate r_{ij} of a link (i, j) depends not only on the power P_{ij} assigned to the link, but also on the interference due to the power assignments on other links. (In this paper we often refer to

P_{ij} as the power assignment, however, the same formulation would clearly apply if P_{ij} represents other types of resource control decisions at link (i, j) , e.g., activation/inactivation, or a random-access attempt-probability.) Let $\vec{P} = [P_{ij}, (i, j) \in \mathcal{L}]$ denote the power assignments and let $\vec{r} = [r_{ij}, (i, j) \in \mathcal{L}]$ denote the data rates. We assume that $\vec{r} = u(\vec{P})$, i.e., the data rates are completely determined by the global power assignment. (One can also extend the model to incorporate channel variations, see Section III-D.) The function $u(\cdot)$ is called the *rate-power function* of the system. There may be constraints on the feasible power assignment. For example, if each node has a total power constraint $P_{i,\max}$, then $\sum_{j:(i,j) \in \mathcal{L}} P_{ij} \leq P_{i,\max}$. Let Π denote the set of feasible power assignments, and let $\mathcal{R} = \{u(\vec{P}), \vec{P} \in \Pi\}$. We assume that $\text{Co}(\mathcal{R})$, the convex hull of \mathcal{R} , is closed and bounded.

There are S users and each user s is associated with a source node f_s and a destination node d_s . Let x_s be the rate with which data is sent from f_s to d_s , over possibly multiple paths and multiple hops. We assume that x_s is bounded in $[0, M_s]$. Each user is associated with a utility function $U_s(x_s)$, which reflects the “utility” to the user s when it can transmit at data rate x_s . We assume that $U_s(\cdot)$ is strictly concave, non-decreasing and continuously differentiable on $[0, M_s]$. The use of such utility functions is common in the congestion control literature to model fairness, since with different utility functions the rate allocations x_s that maximize the total system utility can be mapped to a range of fairness objectives [35], [36].

We assume that time is divided into slots. At each time slot, the scheduling policy will select a power assignment vector \vec{P} (or, equivalently, $\vec{r} = u(\vec{P})$), and select data to be forwarded on each link. Given a user rate vector $\vec{x} = [x_s, s = 1, \dots, S]$, we say that a system is *stable* under a scheduling policy if the queue length at each node remains finite. We can then formulate the following joint congestion-control and scheduling problem:

- **The Congestion-Control Problem:** Find the user rate vector \vec{x} that maximizes the sum of the utilities of all users $\sum_s U_s(x_s)$ subject to the constraint that the system is stable under some scheduling policy.
- **The Scheduling Problem:** For any user rate vector \vec{x} picked by the congestion-control problem, find a scheduling policy that stabilizes the system.

Define the *capacity region* Λ of the system as the largest set of rate vectors \vec{x} such that for any $\vec{x} \in \Lambda$, there exists some scheduling policy that can stabilize the network under the offered-load \vec{x} . Hence, the congestion control part of the problem is simply

$$\begin{aligned} & \max_{x_s \leq M_s} && \sum_s U_s(x_s) && (5) \\ & \text{subject to} && \vec{x} \in \Lambda. \end{aligned}$$

In the sequel, we will review two different ways of stating the capacity region Λ , which then lead to different solutions.

B. The Node-Centric Formulation

In the node-centric formulation, a user rate vector \vec{x} is in the capacity region Λ if and only if there exists a link rate vector \vec{r}^d associated with each destination node d , and the vector $\vec{R} = [\vec{r}^d, d = 1, \dots, N]$ satisfies [15], [37]:

$$\begin{aligned} r_{ij}^d &\geq 0 \text{ for all } (i, j) \in L \text{ and for all } d \\ \sum_{j:(i,j) \in \mathcal{L}} r_{ij}^d - \sum_{j:(j,i) \in \mathcal{L}} r_{ji}^d - \sum_{s:f_s=i, d_s=d} x_s &\geq 0 \\ &\text{for all } d \text{ and for all } i \neq d \end{aligned} \quad (6)$$

$$\left[\sum_d r_{ij}^d \right] \in \text{Co}(\mathcal{R}),$$

where r_{ij}^d can be interpreted as the rate on link (i, j) that is allocated for data towards destination d . These set of equations simply represent a balance of incoming rates and outgoing rates at each node. The convex-hull operator is due to a standard time-averaging argument [15], [37]–[39].

The Solution: Although the rate-power function $u(\cdot)$ is generally a non-convex function, the convex-hull operator in fact makes the capacity region Λ a convex set. Hence, the problem (5) has a dual such that there is no duality gap [28]. Associating a Lagrange multiplier q_i^d for each constraint in (6), we can then obtain the following solution [28], [32], [33]:

- The data rates of the users are determined by

$$x_s(t) = \underset{0 \leq x_s \leq M_s}{\text{argmax}} \left[U_s(x_s) - x_s q_{f_s}^d \right]. \quad (7)$$

- The schedule is determined by first solving the following sub-problem:

$$\vec{r}(t) = \underset{\vec{r} \in u(\vec{P}), \vec{P} \in \Pi}{\text{argmax}} \sum_{(i,j) \in \mathcal{L}} r_{ij} \max_d (q_i^d - q_j^d). \quad (8)$$

Each link then picks the corresponding power assignment that achieves $\vec{r}(t)$, and computes the vectors $\vec{r}^{td}(t)$ as follows: For each link (i, j) , let $d^*(i, j) = \underset{d}{\text{argmax}} (q_i^d - q_j^d)$, and let $r_{ij}^d(t) = r_{ij}(t)$ if $d = d^*(i, j)$ and $r_{ij}^d(t) = 0$, otherwise.

- The Lagrange multipliers are updated by

$$q_i^d(t+1) = \left\{ q_i^d(t) - h_t \left[\sum_{j:(i,j) \in \mathcal{L}} r_{ij}^d(t) - \sum_{j:(j,i) \in \mathcal{L}} r_{ji}^d(t) - \sum_{s:f_s=i, d_s=d} x_s(t) \right] \right\}^+, \quad (9)$$

where $h_t, t = 1, 2, \dots$ is a sequence of positive stepsizes.

The physical interpretation of this set of equations is as follows. The Lagrange multiplier q_i^d can be viewed as a scalar multiple of the queue length at node i for packets destined to node d . Equation (7) corresponds to the congestion control component for determining the data rate of each user. Equation (8) corresponds to the scheduling component. The network first computes the power assignment $\vec{P}(t)$ that corresponds to $\vec{r}(t)$. Then, each link will route data destined to the destination d^* that corresponds to the largest differential backlog $q_i^d - q_j^d$.

Note that given \vec{q} , the congestion control decision and the scheduling decision are made independently. Finally, Equation (9) corresponds to the evolution of the queue length at each node.

One can then show the following convergence result [28].

Proposition 1: If

$$h_t \rightarrow 0 \text{ as } t \rightarrow \infty, \text{ and } \sum_t h_t = \infty,$$

then $\vec{x}(t) \rightarrow \vec{x}^*$ as $t \rightarrow \infty$, where \vec{x}^* is the unique optimal solution to problem (5).

Alternatively, if h_t does not approach zero, then as long as it is small, one can still show that $\vec{x}(t)$ will converge to a small neighborhood around \vec{x}^* . Further, all queues will remain finite, and hence the schedules $\vec{r}^d(t)$ also stabilize the network.

C. Link-Centric Formulation

The link-centric formulation differs from the node-centric formulation in that the capacity constraints are stated as balance equations for each link. For simplicity, we focus on the case where the routes for each user are pre-determined. Let $[H_s^{ij}]$ denote the routing matrix, where $H_s^{ij} = 1$ if traffic of user s passes through link (i, j) , $H_s^{ij} = 0$ otherwise. Then an end-to-end user rate vector \vec{x} belongs to the capacity region Λ if and only if there exists a link rate-vector \vec{r} such that

$$\begin{aligned} \sum_{s=1}^S H_s^{ij} x_s &\leq r_{ij} \text{ for all } (i, j) \in \mathcal{L}, \text{ and} \\ \vec{r} &\in \text{Co}(\mathcal{R}). \end{aligned} \quad (10)$$

The Solution: The capacity region Λ is a convex set, as in the case of the node-centric formulation. Associating a Lagrange multiplier q_{ij} for each constraint in (10), we can obtain the following solution [28], [29]:

- The data rates of the users are determined by

$$x_s(t) = \underset{0 \leq x_s \leq M_s}{\text{argmax}} \left[U_s(x_s) - x_s \sum_{(i,j) \in \mathcal{L}} H_s^{ij} q_{ij} \right]. \quad (11)$$

- The schedule is determined by solving the following subproblem:

$$\vec{r}(t) = \underset{\vec{r} \in u(\vec{P}), \vec{P} \in \Pi}{\text{argmax}} \sum_{(i,j) \in \mathcal{L}} r_{ij} q_{ij}. \quad (12)$$

Each link (i, j) then picks the corresponding power assignment that achieves $\vec{r}(t)$.

- The Lagrange multipliers are updated by

$$q_{ij}(t+1) = \left[q_{ij}(t) + h_t \left(\sum_{s=1}^S H_s^{ij} x_s - r_{ij}(t) \right) \right]^+. \quad (13)$$

where $h_t, t = 1, 2, \dots$ is a sequence of positive stepsizes.

The above solution has a similar physical interpretation to the node-centric solution. Further, a convergence result similar to Proposition 1 can be shown [29]. In the case of cellular networks, the node and link-centric formulations are identical, a case which has been considered in [30].

D. Node-Centric versus Link-Centric Formulation

We make the following remarks:

Decomposition: Given \vec{q} , we have decomposed the problem into a congestion control component ((7) or (11)), and a scheduling component ((8) or (12)), each of which can be solved independently given the Lagrange multipliers q_i^d or q_{ij} . However, the scheduling component of the node-centric formulation incorporates the routing functionality, while the scheduling component of the link-centric formulation does not. Hence, different formulations can result into decomposition between different layering boundaries.

Feedback: In both formulations, given the Lagrange multipliers, each user s then solves its own utility maximization problem (7) or (11) independently. These Lagrange multipliers can be interpreted as the implicit costs for the resources. However, note that in the node-centric formulation, the congestion control component of user s reacts to the implicit cost $q_{f_s}^{d_s}$ at the source node f_s , while in the link-centric formulation, the congestion control component reacts to the sum of the implicit costs along the path of user s . Thus, the node-centric formulation does not require feedback from inside the network, while the link-centric formulation does.

Queues: In both formulations, the Lagrange multipliers may be interpreted as a scalar multiple of the queue length. In the link-centric formulation, each link only needs to maintain one q_{ij} . In the node-centric formulation, each node needs to maintain multiple q_i^d , one per each destination d . Thus, the node-centric formulation has more overhead. On the other hand, in the node-centric formulation, Equation (9) correctly models the way packets move from one node to the other. In contrast, in the link-centric formulation, the Lagrange multiplier q_{ij} will only correspond to the real queue length at link (i, j) under the implicit assumption that the data rate x_s of each user s is applied simultaneously to all links along its path.

Channel variations: Both formulations can be generalized to the case with channel variations (e.g., due to fading and/or mobility of the nodes), in which case only the scheduling component needs to be changed [30], [33], [40].

As a final note, the node-centric formulation can also be generalized to the case with pre-determined routing, and the link-centric formulation can also be generalized to the case with multi-path routing [28]. Primal-dual solutions to this problem are studied in [20], [31].

E. Cases where the Perfect Scheduling Component is Solvable

The scheduling component (8) or (12) is usually difficult to solve because the rate power function $u(\cdot)$ in many wireless settings is not concave. Note that the Lagrange multipliers change every time slot. Hence, a different optimal schedule needs to be computed at every time-slot. Thus, the complexity of the scheduling component is the main bottleneck for the entire solution [41].

A few cases have been studied where computing the exact optimal schedule is of polynomial-time complexity.

- Cellular networks or access-point based single-hop wireless networks where at each time only one node can communicate with the base station or access point [30].
- The so-called *node-exclusive interference model* (i.e., the data rate of each link is fixed at c_{ij} , and the only wireless constraint is that each node can only communicate with one other node at any time), where the optimal schedule corresponds to a Maximum-Weighted-Matching problem [28], [29], [42]. This model is an accurate representation for Bluetooth-like networks and a reasonable approximation to FH-CDMA systems.
- When the data rate of a link is a concave function of *its own power assignment* [43] (which implies no interference between links).

Further, in the so-called *low-SINR* case when the data rate of a link is a linear function of its SINR, the perfect schedule (i.e., power assignment) satisfies the property that each node either does not transmit at all, or transmits at the maximum power to only one other node. This property significantly reduces the search space for optimal power assignments [28], [38]. However, the complexity remains exponential in the number of nodes.

For certain special rate-power functions, although the function $u(\vec{P})$ is originally not a concave function of \vec{P} , it may become a concave function after some change of variables. Such cases include: (a) the so-called *high-SINR* case where the rate r_{ij} of each link (i, j) is a logarithmic function of its SINR [44], [45]; (b) the low-SINR case [46]; (c) single-channel Aloha networks [47]–[49]. With these rate-power functions, rather than solving the problem (5), several researchers instead study the following problem [44]–[49] (the link-centric formulation is used here):

$$\begin{aligned} & \max_{x_s \leq M_s} \sum_{s=1}^S U_s(x_s) \\ & \text{subject to} \quad \sum_{s=1}^S H_s^{ij} x_s \leq r_{ij}, \text{ for all } (i, j) \in \mathcal{L}, \text{ and} \\ & \quad \vec{r} \in \mathcal{R}. \end{aligned} \tag{14}$$

Note the omission of the convex-hull operator compared with (10). Recall that $\mathcal{R} = \{\vec{r}' | \vec{r}' = u(\vec{P}), \vec{P} \in \Pi\}$. If $u(\cdot)$ is not concave, then \mathcal{R} is not a convex set. However, for the special cases listed above, by some change of variables $\vec{r}' \rightarrow \vec{r}'$, $\vec{P} \rightarrow \vec{P}'$, it is possible to rewrite the problem (14) as a convex program of \vec{r}' and \vec{P}' (sometimes requiring additional assumptions on the utility function $U_s(\cdot)$). One can then use standard convex programming techniques to solve problem (14). While such transformations to convex programs typically make the problem much easier to solve, one should be careful about this approach in general since it may incur a loss of capacity due to the omission of the convex-hull operator in (14). (The omission of the convex hull does not allow time-interleaving of schedules). Finally, the above transformation approach critically depends on the rate-power function used, and may not be applicable for other interference models.

IV. USING IMPERFECT SCHEDULING IN THE CROSS-LAYER SOLUTION

Apparently, the complexity of the scheduling component has become the main bottleneck for the cross-layer solution. It would be even harder to come up with distributed solutions that are suitable for large networks. A natural question is then: what if we relax the optimality requirement of the scheduling component ((8) or (12)), and let the scheduling component only compute an suboptimal schedule at each time. As we will see in Section V, this relaxation opens up a large class of new scheduling solutions that are simple, provably efficient, and can be implemented in a distributed fashion. However, we first need to understand how to use these imperfect scheduling policies in the cross-layer solution. Recall that the perfect schedules computed by the scheduling component (8) and (12) are in fact the throughput-optimal schedules that achieve the largest capacity region Λ . An imperfect scheduling policy will typically achieve a smaller capacity region. We next review two approaches to use such imperfect scheduling policies for solving the joint congestion-control and scheduling problem.

A. The Layered Approach

If the imperfect scheduling policy can ensure a rate region $\Lambda_0 \subset \Lambda$, and Λ_0 can be described by a set of simple linear constraints, then we can use the following *layered approach*. We can compute the end-to-end user rate vector \vec{x} by solving the congestion control problem (5) with the constraints replaced by Λ_0 . Since now the “optimal” rate allocation lies in Λ_0 , the imperfect scheduling policy will then be able to support the computed end-to-end user rate-vector.

In [50], such a layered approach is applied to the node-exclusive interference model. Under this model, a simple scheduling policy called *Maximal Matching* can guarantee the following capacity region Λ_0 (see Section V-A): the user rate-vector \vec{x} belongs to Λ_0 if and only if

$$\sum_{(i,j) \in \mathcal{L}} \frac{\sum_{s=1}^S H_s^{ij} x_s}{c_{ij}} + \sum_{(j,i) \in \mathcal{L}} \frac{\sum_{s=1}^S H_s^{ji} x_s}{c_{ji}} \leq \frac{1}{2}, \text{ for all node } i. \quad (15)$$

It will be shown in Section V-A that $\Lambda_0 \supset \Lambda/2$. Hence, the loss of efficiency is at most 1/2. The constraint in Λ_0 is essentially the same as in wireline networks. We can then associate Lagrange multipliers for each constraint in (15), and solve the optimal user rate-vector \vec{x} in Λ_0 that maximizes the total system utility. The resulting user rate-vector \vec{x} can also be stabilized by the Maximal Matching policy.

B. The Cross-Layered Approach

In the cross-layered approach, we still use the cross-layer solutions in Section III, except that we replace the scheduling component (8) or (12) by an imperfect scheduling policy. Hence, the dynamics of the congestion control component and the scheduling component are still coupled through the implicit costs (i.e., Lagrange multipliers). Note that once the scheduling component of the optimal solution is replaced by

an imperfect scheduling policy, the dynamics of the system can become much more complicated. The concern here is that, due to the interaction between the two components, the system may get stuck at a local optimum which may be globally sub-optimal and quite inefficient. The following questions need to be addressed: (a) Does the system still converge to some user rate-vector? (b) Is the resulting user rate-vector still fair and efficient?

These questions are addressed in [29] for a class of imperfect scheduling policies called S_γ -policies. An S_γ -policy will compute, at each time t , a schedule $\vec{r}(t)$ that satisfies:

$$\sum_{l=1}^L r_l(t) q^l(t) \geq \gamma \max_{\vec{r} \in \mathcal{R}} \sum_{l=1}^L r_l q^l(t), \quad (16)$$

where γ is a constant in $(0, 1]$. It is known that an S_γ -policy can guarantee a capacity region of at least $\gamma\Lambda$. Thus, the parameter γ can be viewed as a tuning parameter indicating the degree of imprecision of the imperfect schedule. A number of low-complexity scheduling algorithms fall into this class. For example, under the node-exclusive interference model, simple Greedy Maximal Matching policy [29], [51] is an S_γ -policy with $\gamma = 1/2$.

In [29], the fairness and efficiency of a cross-layer solution using an S_γ -policy are studied. Two types of results are reported.

- In the *static* case where the user population is fixed, [29] shows that, if the cross-layer solution with an S_γ -policy converges to some rate-allocation, the rate-allocation satisfies a weak fairness property. In particular, if the utility function is logarithmic (i.e., of the form $U_s(x_s) = w_s \log x_s$, where w_s is the weight of user s), then the convergent point \vec{x} will satisfy

$$\vec{x} \in \Lambda \text{ and } \sum_{s=1}^S w_s \frac{\gamma x_s^*}{x_s} \leq \sum_{s=1}^S w_s, \quad (17)$$

where x_s^* is the optimal rate-allocation. In other words, the rate of each user is unlikely to be too unfair compared to γx_s^* . However, in general it is impossible to establish the convergence of \vec{x} to any particular point.

- In the *dynamic* case, i.e., in which the users enter and leave the system according to a stochastic process, [29] establishes the lower bound on the stability region of the system employing the cross-layer solution with imperfect scheduling. In particular, assuming that users of class s arrive according to a Poisson process with rate λ_s and that each user brings with it a file for transfer whose size is exponentially distributed with mean $1/\mu_s$, we can define $\rho_s = \lambda_s/\mu_s$ as the load brought by users of class s . Let $\vec{\rho} = [\rho_1, \dots, \rho_S]$. Then [29] shows that the stability region of the system (i.e., the set of offered-load vectors $\vec{\rho}$ under which the system remains stable) is at least $\gamma\Lambda$ when we use the cross-layer solution with an S_γ -policy.

The above results tell us that although only a weak fairness property can be shown for the static case, it appears to be sufficient to ensure a reasonable stability region. Thus, the

potential local sub-optima that could occur in the static case appears to be inconsequential when the arrivals and departures of the users are taken into account.

Finally, we would like to point out that not all imperfect scheduling policies are of the S_γ -type. How to address the cross-layer impact of these other imperfect scheduling policies is still an open problem.

V. DISTRIBUTED SCHEDULING ALGORITHMS

We next review a class of imperfect scheduling algorithms. We are interested in these algorithms because they are very simple and can be implemented in a distributed fashion. We will focus on showing that these algorithms can also guarantee a certain fraction of the overall capacity region. These scheduling algorithms can then be used to solve the joint congestion-control and scheduling problem, either via the layered approach in Section IV-A, or via the cross-layer approach in Section IV-B [29].

We consider a fairly general interference model for the wireless network. There are L links operating in a discrete time-slotted manner. Each link can send c_l packets per time-slot, provided that there are no other interfering links activated at the same time; otherwise, if two links that interfere with each other are activated at the same time, none of the links can send any packets. Associated with each link is a set of other links which interfere with the given link. The only assumption that we make on the interference relationship is that it is symmetric, i.e., if link l interferes with link k , then l also interferes with k . Specifically, let \mathcal{E}_l be the interference set associated with link l , then: (a) $l \in \mathcal{E}_l$, and (b) if $k \in \mathcal{E}_l$, then $l \in \mathcal{E}_k$.

We are interested in maximal policies that select links for scheduling such that the interference constraints are satisfied. The policies are maximal in the sense that the schedule can be chosen in the following fashion: start with any link in the network and add links to the schedule, one link at a time, subject to interference constraints. If the scheduler proceeds to implement this algorithm until no more links can be added to the schedule, then it is called maximal. Such maximal policies require no significant computation, but may require an overhead for implementation in a decentralized manner as we will discuss later.

It turns out that for the purpose of proving capacity regions, we can focus on the single-hop case, i.e., by assuming that each user's route consists of a link. Then under certain assumptions, these results can be extended to the case with multi-hop routes⁴. Associated with each link is a stochastic arrival process $\{A_l(n)\}$, where $A_l(n)$ is the number of packet arrivals in slot n to link l . We assume that the arrival processes are stationary and let $\lambda_l = E(A_l(n))$. The number of departures from link l at time n is denoted by $D_l(n)$, and $D_l(n)$ is

assumed to be less than or equal to c_l packets. Let $q_l(n)$ be the number of packets at link l at the beginning of time slot n . Let $d_l(n)$ be an indicator function indicating whether link l is scheduled in time slot n or not. We make the assumption that link l is eligible for scheduling only if it has at least c_l waiting packets. Thus, $D_l(n) = c_l d_l(n)$. We assume that the sequence of events in each time slot is as follows: (i) schedule consistent with the interference constraint is chosen, (ii) departures occur next, and (iii) arrivals occur last.

We now formally define a maximal scheduling policy. A policy is said to be a maximal policy if the departures under this policy satisfy the following constraint: if $q_l(n) \geq c_l$, then $\sum_{k \in \mathcal{E}_l} d_k(n) = 1$. This ensures that when a link has at least c_l packets, either the link is scheduled or if it is not scheduled, then the reason that it is not scheduled is that another link from its interference set is scheduled. The above policy is a natural extension of the maximal schedules considered for high-speed switches in [53], [54] and, for Bluetooth-like wireless networks, in [28], [29].

The simplest model to consider is one where

$$A(n) := \{A_1(n), A_2(n), \dots, A_L(n)\}$$

are *i.i.d.* across n . In other words, the arrival process is *i.i.d.* across time slots, but may be dependent across links. The extension to more general Markovian arrival processes is straightforward. Define the state of the system to be

$$Q(n) := (q_1(n), q_2(n), \dots, q_L(n)),$$

where the dynamics of $q_l(n)$ are given by

$$q_l(n+1) = q_l(n) + A_l(n) - D_l(n).$$

Assume that $\{D_l(n)\}$ is chosen according to some probability distribution given $Q(n)$, i.e., $P(\{D_l(n)\}|Q(n))$ is given. Thus, $Q(n)$ is a countable-state-space Markov chain. Note that $P(\{D_l(n)\}|Q(n))$ can be arbitrary. In particular, $\{D_l(n)\}$ could be any sequence of schedules consistent with the interference constraints. In a real network, the maximal schedule may be computed by a random access protocol. In this case, it may be reasonable to assume that all feasible maximal schedules are equally likely at each time instant. However, the result applies to more general models as well; in fact, it holds for any rule used to choose the set of active links, as long as the resulting schedule is a maximal schedule. The following result was proved in [55], [56].

Proposition 2: For any distributions $P(\{D_l(n)\}|Q(n))$, the Markov chain $Q(n)$ is stable-in-the-mean, i.e.,

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \sum_{l=1}^L E(q_l(n)) < \infty,$$

if

$$\sum_{k \in \mathcal{E}_l} \frac{\lambda_k}{c_k} < 1, \quad \forall l. \quad (18)$$

Remark: The condition in (18) can be used to form the region Λ_0 (see Section IV-A). One may then use the layered approach in Section IV-A to solve the joint congestion-control and scheduling problem.

⁴For example, if we assume that the traffic on a route arrives at all the links on a route instantaneously, then it is effectively the same problem as one where the routes consist of a single hop. Alternatively, it has been shown in [52] that if one uses traffic re-shaping at each link, the stability conditions for single-hop routes also apply to multi-hop routes.

A. Worst-Case Loss in Throughput for Node-Exclusive Interference Models

The simplest interference model, motivated by FH-CDMA and Bluetooth networks, is one where the only constraint on a node's transceiver is that it cannot receive and transmit simultaneously. This is called the *node-exclusive interference model*. In graph-theoretic terms, this means a valid schedule in such a network can be represented as a matching. A *matching* is a collection of links in the network such that no two edges have a common node. If no two links chosen to be part of the schedule share a common node, then it means that no node transmits or receives simultaneously, thus satisfying the node-exclusive interference model constraint. The maximal schedule mentioned previously in this section can then be thought of as a maximal matching. (A *maximal matching* is a matching to which no link can be added without violating the *matching* property.) The specialization of the result in Proposition 2 to the node-exclusive interference model was first obtained in [29] for a more general traffic model where the number of users in the network was also allowed to vary according to a stochastic process. In this subsection, we discuss the loss in throughput due to maximal matching schedules. In the rest of this section, we will assume that $c_l = 1 \forall l$.

Now, consider any link l and its interference set \mathcal{E}_l . The interference set \mathcal{E}_l consists of all links that have a node in common with link l . A little thought shows that a maximum of two links in \mathcal{E}_l can be scheduled: one link connected to each of the nodes associated with l . This observation immediately suggests that any set of arrival rates that can be supported in the interference set \mathcal{E}_l must satisfy

$$\sum_{k \in \mathcal{E}_l} \lambda_k < 2. \quad (19)$$

The above result follows from simple queueing-theoretic considerations: the maximum number of packets that can be drained from this interference set at any time instant is two, and thus the arrival rate has to be less than 2 for stability. Earlier in this section, we proved that any arrival rate such that $\sum_{k \in \mathcal{E}_l} \lambda_k < 1$ can be scheduled using a maximal schedule (i.e., a maximal matching under the node-exclusive model). Thus, the maximum loss in throughput under the node-exclusive interference model is $1/2$.

Under the node-exclusive model, the maximal matching policy may also be viewed as an S_γ -policy with respect to a system where the capacity constraint is given by (19). One may then use the cross-layer approach of Section IV-B to construct fully distributed solutions to the joint congestion-control and scheduling problem. See [40] for details.

B. Worst-Case Loss in Throughput for General Interference Models

In the previous subsection, we showed that, under the node-exclusive interference model, the worst-case throughput loss is bounded by a factor $1/2$ independent of the topology of the network. However, such a topology-independent result appears to be difficult to obtain for general interference models. For

general interference models, suppose that κ_l be the maximum number of links of \mathcal{E}_l that can be scheduled simultaneously. Then, as before, an elementary queueing argument shows that the set of stabilizable arrival rates should satisfy

$$\sum_{k \in \mathcal{E}_l} \lambda_k < \kappa_l.$$

Letting $\kappa_{max} = \max_l \kappa_l$, we see that the set of stabilizable arrival rates must satisfy

$$\sum_{k \in \mathcal{E}_l} \frac{\lambda_k}{\kappa_{max}} < 1.$$

Comparing this to the sufficient condition for stability under maximal schedules, we see that the throughput is reduced by a factor of at most $1/\kappa_{max}$. In [57], it has been shown that the above bound for throughput loss is tight for certain types of arrival processes and topologies. However, the argument requires that the arrival processes to the various links in the network are correlated.

C. Implementation Considerations and a Constant-Overhead Algorithm

Maximal schedules can be implemented in a distributed fashion as follows. Suppose each node uses a small packet to request a transmission to a neighboring node (such as a Request-to-Transmit, or RTS, packet in 802.11). If the receiving node has not already scheduled to transmit, and it has not received RTS from some other node, it can respond with a Clear-to-Send (CTS) packet, and the sending node and the receiving node are then matched. Otherwise, if a CTS is not received by the sending node, the sending node could then try to connect to another node. At the end of this process, one would have a maximal schedule. While this is a distributed algorithm, it comes at the cost of a high overhead. Note that it may take many RTS/CTS cycles for a maximal schedule to be established while the results of Proposition 2 assumes that this process is instantaneous. Thus, there could be a further loss in throughput beyond what the analysis indicates due to the overhead.

The overhead due to the RTS/CTS implementation can be reduced under some ideal time-synchronization assumptions. Assume that a small portion of time at the beginning of each time-slot is allocated to compute the schedule. Suppose that this control time-slot has a duration T . Then, at the beginning of each control time-slot, each backlogged link generates a uniform random number in $[0, T]$. Note that each *link* generates a random number; it is not sufficient for each *node* to generate a random number. Each link then sends out an RTS at the time indicated by its random number unless it has previously heard an RTS. In this manner, all collisions are avoided and a maximal schedule is found by time T . If T is small compared to the duration of the entire time-slot, then the overhead is small. While time synchronization may be a reasonable assumption, it is unreasonable to assume that the RTS can be sent in an arbitrarily small amount of time. In reality, it takes some time to transmit and receive an

RTS. Thus, the control time-slot duration has to be thought as a finite number of discrete mini-slots. In particular, say M such mini-slots can be fitted into the time interval T . Then, each node has to pick a random number in $\{1, 2, \dots, M\}$. Since it is now a discrete random variable, collisions can happen. To reduce the probability of collisions, one may have to choose M sufficiently large which may again lead to a significant overhead. In particular the choice of M will depend upon the number of nodes in the network and thus, the solution is not scalable. These considerations then motivate the question of whether there exists a *constant-overhead* algorithm to implement a distributed schedule. By *constant-overhead*, we mean the overhead is independent of the network size or topology. Recently, it has been shown that indeed a constant-overhead algorithm does exist provided that we are willing to tolerate a further small reduction in throughput beyond the loss in throughput dictated by the maximal schedule [58]. The algorithm is similar in spirit to an algorithm suggested earlier for switches [51].

VI. CONCLUSIONS AND OPEN PROBLEMS

In the previous sections, we have summarized the significant progress in optimization-based cross-layer resource allocation in wireless networks. In particular, we note that the formulation can be viewed as an approach to clean-slate design for communication networks in general, wireline or wireless networks. A noteworthy feature of this approach is that the algorithms obtained from this clean-slate design naturally map to different layers of the protocol stack: congestion control at the transport layer, routing at the network layer, and scheduling/power control at the MAC/PHY layer. A key difference between existing network architectures and the optimization approach is that the optimization approach also clearly delineates the interactions necessary between the various layers to achieve optimal performance. While the progress in this area has been substantial over the last few years, there are still several open problems, especially in the area of multi-hop wireless networks, where scalability becomes a critical issue.

Propagation delays and connection-level stability analysis: While the focus of much of this paper is on resource allocation in networks with a fixed number of users, we had discussed connection arrivals and departures in the context of imperfect schedules (in Section IV). For the case of perfect schedules, it was proven in [59] that the class of joint congestion-control/routing/MAC policies discussed in Section III is throughput optimal when there are connection-level arrivals and departures. However, it was assumed that there were no propagation delays between the nodes in the network. While this may be a reasonable assumption in small wireless networks, it is not a reasonable assumption when one considers heterogeneous networks, where part of the network may be a wired network such as the Internet. Thus, an open issue is to prove the throughput optimality of the algorithms in the presence of propagation delays.

Traffic arriving at all nodes on the route instantaneously: In Section III, the node-centric formulation of the resource allocation problem in multi-hop wireless networks was able to model packets traversing the nodes in the network one at a time. However, in the link-centric formulation, we have assumed that routes were pre-computed and a packet arrives at all the links on its path simultaneously. We proved stability and convergence for the link-centric solution under this assumption. It is well-known in queue-theoretic literature that queue-length stability when packets are assumed to arrive at all nodes on a route simultaneously does not guarantee the stability of networks where the packets have to traverse one node at a time [60], [61]. One solution to this problem was suggested in [50], [52] using a policing device called a *regulator* to re-shape traffic at each node in the network. However, this solution requires the use of per-flow queueing in the network. On the other hand, since the networks in this paper are congestion-controlled and are unlike the examples in [60], [61], it is unclear that the type of instability exhibited in [60], [61] apply to the models considered in this paper. Resolving this question and proving stability in general is an open problem.

Complexity issues and distributed algorithms for general physical interference models: While the general resource allocation framework in Sections III and IV allows for arbitrary physical layer model, the maximal schedules and other distributed algorithms of Section V focused only on simple collision models. While it is interesting to obtain algorithms that achieve a guaranteed fraction of the maximum possible throughput for these practical interference models of Section V, it is also true that emerging networks have more sophisticated mechanisms available to control interference, such as power control and the use of multiple carriers. Designing algorithms with guaranteed minimum throughput levels under general interference models is indeed a challenging open problem.

Cross-layer design with fairness: In Section IV, we noted that there are two methods to make use of the reduced-throughput scheduling algorithms of Section V (such as maximal schedules) along with the transport-layer congestion-control algorithms. One way is to treat the congestion signals as though they are emanating from an algorithm that achieves full throughput. In this approach, which we called the *cross-layer approach*, strict fairness is not guaranteed and indeed the congestion control algorithm with a fixed number of users may not converge. However, at the connection level, the number of active users/connections in the system remains stochastically bounded provided that the load imposed on the system by connection arrivals and departures lies within the guaranteed stability region. Further, the cross-layer approach is typically able to exploit the excess bandwidth in the sense that the actual stability region is usually much larger than the worst-case stability region guaranteed by the reduced-throughput scheduling algorithm. A second approach is to impose the minimum guaranteed throughput as a resource constraint and use this constraint to generate congestion signals. This approach, which we called the *layered approach*,

guarantees fairness in the reduced rate region but may not be able to exploit excess bandwidth. An interesting open question is whether there exists a protocol that combines the desirable features of the two algorithms, i.e., one that can guarantee the minimum bandwidths dictated by the fair-shares in the reduced rate region while being able to exploit any excess bandwidth beyond the minimum guaranteed bandwidth.

Constant-overhead/low-overhead implementation: In Section V-C, we briefly discussed the need for contention resolution protocols that assure a constant overhead independent of network topology and size. One such constant-overhead distributed scheduling algorithm was proposed in [58]. However, its guaranteed throughput is lower than the guaranteed throughput of maximal schedules of Section V. Recently, distributed algorithms that achieve full throughput have been discussed in [62] and [63]. However, these algorithms require significant overheads or require unacceptably large delays. An open problem is to characterize the tradeoff between the overhead required and the reduction in throughput, if any, due to the requirement of low overhead in the contention resolution algorithm. The ultimate goal would be to develop low-overhead distributed algorithms that achieve the maximum possible throughput in the network.

Tightness of throughput-loss bounds: All the results on bounding the throughput loss for imperfect schedules are based on worst-case analysis. In other words, all the results so far only give lower bounds on the achievable throughput. However, simulations in [29] suggest that the average throughput (over all imperfect schedules) could be much larger. An interesting open issue is to characterize the average user-perceived throughput using imperfect schedules.

Non-Concave utility functions: Finally, all of the developed results have assumed that the utility functions are concave functions of the rate. These utility functions are appropriate for highly elastic traffic, such as data traffic. However, for other applications (e.g., voice or video), a more appropriate characterization of the user utility would be via non-concave sigmoidal utility functions [64]–[66]. An open problem is to develop solutions to the cross-layer resource allocation problem in this setting.

REFERENCES

- [1] F. P. Kelly, A. Maulloo, and D. Tan, "Rate Control in Communication Networks: Shadow Prices, Proportional Fairness and Stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [2] H. Yaiche, R. Mazumdar, and C. Rosenberg, "A Game Theoretic Framework for Bandwidth Allocation and Pricing in Broadband Networks," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 667–678, Oct. 2000.
- [3] S. H. Low and D. E. Lapsley, "Optimization Flow Control—I: Basic Algorithm and Convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, December 1999.
- [4] S. Kunniyur and R. Srikant, "End-to-End Congestion Control Schemes: Utility Functions, Random Losses and ECN Marks," in *Proceedings of IEEE INFOCOM*, Tel-Aviv, Israel, March 2000.
- [5] S. H. Low and R. Srikant, "A Mathematical Framework for Designing a Low-Loss Low-Delay Internet," *Network and Spatial Economics*, vol. 4, no. 1, pp. 75–102, March 2004.
- [6] R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.
- [7] M. Chiang, S. H. Low, R. A. Calderbank, and J. C. Doyle, "Layering as Optimization Decomposition," *Proceedings of IEEE*, December 2006.
- [8] R. Knopp and P. Humblet, "Information capacity and power control in single-cell multiuser communications," in *Proceedings of the ICC*, 1995.
- [9] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi, "CDMA/HDR: a bandwidth-efficient high-speed wireless data service for nomadic users," *IEEE Communications magazine*, pp. 70–77, July 2000.
- [10] S. Borst and P. Whiting, "Dynamic rate control algorithms for CDMA throughput optimization," in *Proceedings of IEEE Infocom 2001*, Alaska, April 2001.
- [11] —, "The use of diversity antennas in high-speed wireless systems: Capacity gains, fairness issues, multi-user scheduling," *Bell Laboratories Technical Memorandum*, 2001.
- [12] X. Liu, E. K. P. Chong, and N. B. Shroff, "Opportunistic transmission scheduling with resource-sharing constraints in wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 19, no. 10, October 2001.
- [13] —, "A Framework for Opportunistic Scheduling in Wireless Networks," *Computer Networks*, vol. 41, no. 4, pp. 451–474, March 2003.
- [14] A. Jalali, R. Padovani, and R. Pankaj, "Data throughput of CDMA-HDR a high efficiency-high data rate personal communication wireless system," in *IEEE VTC*, 2000.
- [15] L. Tassiulas and A. Ephremides, "Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multihop Radio Networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936–1948, December 1992.
- [16] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, P. Whiting, and R. Vijayakumar, "Providing quality of service over a shared wireless link," *IEEE Communications Magazine*, vol. 39, pp. 150–153, February 2001.
- [17] S. Shakkottai and A. Stolyar, "Scheduling for multiple flows sharing a time-varying channel: The exponential rule," *Translations of the AMS*, 2001, a volume in memory of F. Karpelevich.
- [18] —, "Scheduling of a shared a time-varying channel: The exponential rule stability," in *INFORMS Applied Probability Conference*, New York, July 2001.
- [19] R. Buche and H. J. Kushner, "Control of mobile communication systems with time-varying channels via stability methods," *IEEE Transactions on Automatic Control*, 2004.
- [20] A. L. Stolyar, "Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm," *Queueing Systems*, vol. 50, no. 4, pp. 401–457, 2005.
- [21] S. Borst, "User-level performance of channel-aware scheduling algorithms in wireless data networks," *IEEE/ACM Transactions on Networking*, vol. 13, no. 3, 2005.
- [22] S. Shakkottai, "Effective capacity and QoS for wireless scheduling," 2004, preprint.
- [23] L. Ying, R. Srikant, A. Eryilmaz, and G. Dullerud, "A large deviations analysis of scheduling in wireless networks," 2005, submitted. Earlier versions of the paper appeared in the IEEE CDC 2004, IEEE CDC 2005 and IEEE ISIT 2006.
- [24] S. Sarkar and L. Tassiulas, "End-to-end Bandwidth Guarantees Through Fair Local Spectrum Share in Wireless Ad-hoc Networks," in *Proceedings of the IEEE Conference on Decision and Control*, Maui, Hawaii, December 2003.
- [25] Y. Yi and S. Shakkottai, "Hop-by-hop Congestion Control over a Wireless Multi-hop Network," in *Proceedings of IEEE INFOCOM*, Hong Kong, March 2004.
- [26] Y. Xue, B. Li, and K. Nahrstedt, "Price-based Resource Allocation in Wireless Ad Hoc Networks," in *Proceedings of the Eleventh International Workshop on Quality of Service (IWQoS 2003)*, also *Lecture Notes in Computer Science*, ACM Springer-Verlag, vol. 2707, Monterey, CA, June 2003, pp. 79–96.
- [27] L. Chen, S. H. Low, and J. C. Doyle, "Joint Congestion Control and Media Access Control Design for Wireless Ad Hoc Networks," in *Proceedings of IEEE INFOCOM*, Miami, FL, March 2005.
- [28] X. Lin and N. B. Shroff, "Joint Rate Control and Scheduling in Multihop Wireless Networks," in *Proceedings of the IEEE Conference on Decision and Control*, Paradise Island, Bahamas, December 2004.
- [29] —, "The Impact of Imperfect Scheduling on Cross-Layer Rate Control in Multihop Wireless Networks," in *Proceedings of IEEE INFOCOM*, Miami, FL, March 2005.

- [30] A. Eryilmaz and R. Srikant, "Fair Resource Allocation in Wireless Networks Using Queue-length-based Scheduling and Congestion Control," in *Proceedings of IEEE INFOCOM*, Miami, FL, March 2005.
- [31] —, "Joint Congestion Control, Routing and MAC for Stability and Fairness in Wireless Networks," in *Proc. International Zurich Seminar on Communications*, 2006, to appear.
- [32] I. Paschalidis, W. Lai, and D. Starobinski, "Asymptotically Optimal Transmission Policies for Low-Power Wireless Sensor Networks," in *Proceedings of IEEE INFOCOM*, Miami, FL, March 2005.
- [33] M. J. Neely, E. Modiano, and C. Li, "Fairness and Optimal Stochastic Control for Heterogeneous Networks," in *Proceedings of IEEE INFOCOM*, Miami, FL, March 2005.
- [34] M. Johansson and L. Xiao, "Scheduling, Routing and Power Allocation for Fairness in Wireless Networks," in *IEEE VTC - Spring*, Milan, Italy, May 2004.
- [35] T. Bonald and L. Massoulié, "Impact of Fairness on Internet Performance," in *Proceedings of ACM Sigmetrics*, Cambridge, MA, June 2001, pp. 82–91.
- [36] J. Mo and J. Walrand, "Fair End-to-End Window-Based Congestion Control," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 556–567, October 2000.
- [37] M. J. Neely, E. Modiano, and C. E. Rohrs, "Dynamic Power Allocation and Routing for Time Varying Wireless Networks," in *Proceedings of IEEE INFOCOM*, San Francisco, April 2003.
- [38] R. L. Cruz and A. V. Santhanam, "Optimal Routing, Link Scheduling and Power Control in Multi-hop Wireless Networks," in *Proceedings of IEEE INFOCOM*, San Francisco, April 2003.
- [39] S. Toumpis and A. J. Goldsmith, "Capacity Regions for Wireless Ad Hoc Networks," *IEEE Transactions on Wireless Communications*, vol. 2, no. 4, pp. 736–748, July 2003.
- [40] X. Lin and N. B. Shroff, "The Impact of Imperfect Scheduling on Cross-Layer Congestion Control in Wireless Networks," *IEEE/ACM Transactions on Networking*, to appear.
- [41] G. Sharma, R. R. Mazumdar, and N. B. Shroff, "Maximum Weighted Matching with Interference Constraints," in *IEEE International Workshop on Foundations and Algorithms For Wireless Networking*, Pisa, Italy, March 2006.
- [42] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Englewood Cliffs, New Jersey: Prentice-Hall, 1982.
- [43] L. Xiao, M. Johansson, and S. Boyd, "Simultaneous Routing and Resource Allocation via Dual Decomposition," in *Proceedings of 4th Asian Control Conference*, Singapore, September 2002, pp. 29–34.
- [44] M. Chiang, "Balancing Transport and Physical Layer in Multihop Wireless Networks: Jointly Optimal Congestion and Power Control," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 1, pp. 104–116, January 2005.
- [45] —, "Geometric Programming for Communication Systems," *Foundations and Trends in Communications and Information Theory*, vol. 2, no. 1-2, pp. 1–154, July 2005.
- [46] M. Arisoylu, T. Javidi, and R. L. Cruz, "End-to-End and Mac-Layer Fair Rate Assignment in Interference Limited Wireless Access Networks," in *Proceedings of IEEE ICC 2006*, Istanbul, Turkey, June 2006.
- [47] X. Wang and K. Kar, "Cross-Layer Rate Control for End-to-End Proportional Fairness in Wireless Networks with Random Access," in *ACM Mobihoc*, Urbana-Champaign, IL, May 2005.
- [48] J. W. Lee, M. Chiang, and R. A. Calderbank, "Jointly Optimal Congestion and Contention Control in Wireless Ad Hoc Networks," *IEEE Communications Letters*, vol. 10, no. 3, pp. 216–218, March 2006.
- [49] J. Zhang and D. Zheng, "A Stochastic Primal-Dual Algorithm for Joint Flow Control and MAC Design in Multi-hop Wireless Networks," in *Conference on Information Sciences and Systems*, Princeton, New Jersey, March 2006.
- [50] L. Bui, A. Eryilmaz, R. Srikant, and X. Wu, "Joint congestion control and distributed scheduling in multihop wireless networks with a node-exclusive interference model," in *Proceedings of IEEE INFOCOM*, 2006.
- [51] E. Leonardi, M. Mellia, F. Neri, and M. A. Marsan, "On the Stability of Input-Queued Switches with Speed-Up," *IEEE/ACM Transactions on Networking*, vol. 9, no. 1, pp. 104–118, February 2001.
- [52] X. Wu and R. Srikant, "Regulated maximal matching: A distributed scheduling algorithm for multi-hop wireless networks with node-exclusive spectrum sharing," in *Proceedings of the IEEE Conference on Decision and Control*, 2005.
- [53] T. Weller and B. Hajek, "Scheduling non-uniform traffic in a packet-switching system with small propagation delay," *IEEE/ACM Transactions on Networking*, pp. 813–823, December 1997.
- [54] J. Dai and B. Prabhakar, "The throughput of data switches with and without speedup," in *Proceedings of IEEE INFOCOM*, 2000, pp. 556–564.
- [55] X. Wu, R. Srikant, and J. R. Perkins, "Queue-Length Stability of Maximal Greedy Schedules in Wireless Network," in *Proceedings of Information Theory and Applications Inaugural Workshop*, University of California, San Diego, February 2006.
- [56] P. Chaporkar, K. Kar, and S. Sarkar, "Achieving Queue Length Stability Through Maximal Scheduling in Wireless Networks," in *Proceedings of Information Theory and Applications Inaugural Workshop*, University of California, San Diego, February 2006.
- [57] —, "Throughput Guarantees in Maximal Scheduling in Wireless Networks," in *Proceedings of 43d Annual Allerton Conference on Communication, Control and Computing*, Monticello, IL, September 2005.
- [58] X. Lin and S. Rasool, "Constant-Time Distributed Scheduling Policies for Ad Hoc Wireless Networks," *Submitted, also available at <http://min.ecn.purdue.edu/~linx/papers.html>*, 2006.
- [59] X. Lin, N. Shroff, and R. Srikant, "On the connection-level stability of congestion-controlled communication networks," 2006, Technical Report.
- [60] P. R. Kumar and T. I. Seidman, "Dynamic instabilities and stabilization methods in distributed real-time scheduling of manufacturing systems," *IEEE Transactions on Automatic Control*, pp. 289–298, March 1990.
- [61] A. Rybko and A. Stolyar, "Ergodicity of Stochastic Processes Describing the Operation of Open Queueing Networks," *Problems of Information Transmission*, vol. 28, 1992, translated from Problemy Peredachi Informatsii, vol. 28, no. 3, pp. 3–26, 1992.
- [62] A. Eryilmaz, E. Modiano, and A. Ozdaglar, "Distributed control for throughput-optimality and fairness in wireless networks," preprint, 2006.
- [63] Y. Yi, G. de Veciana, and S. Shakkottai, "Learning contention patterns and adapting to load/topology changes in a mac scheduling algorithm," preprint, 2006.
- [64] J.-W. Lee, R. R. Mazumdar, and N. B. Shroff, "Downlink power allocation for multi-class CDMA wireless networks," in *IEEE Infocom'02*, vol. 3, 2002, pp. 1480–1489.
- [65] J. W. Lee, R. R. Mazumdar, and N. B. Shroff, "Opportunistic power scheduling for multi-server wireless systems with minimum performance constraints," in *IEEE INFOCOM 2004*, Hong Kong, China, 2004.
- [66] M. Chiang, S. Zhang, and P. Hande, "Distributed rate allocation for inelastic flows: Optimization frameworks, optimality conditions, and optimal algorithms," in *IEEE INFOCOM '05, Miami, Florida, Mar. 2005*.



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

Dr. Lin's research interests are resource allocation, optimization, network pricing, routing, congestion control, network as a large system, cross-layer design in wireless networks, mobile ad hoc and sensor

networks. He received the 2005 best paper of the year award from *Journal of Communications and Networks*. His paper was also one of two runner-up papers for the best-paper award at IEEE INFOCOM 2005.

The biography for Ness B. Shroff and R. Srikant can be found on page xxx and page xxx of this issue.