

Optimal Pricing in a Free Market Wireless Network

Michael J. Neely

University of Southern California

<http://www-rcf.usc.edu/~mjneely>

Abstract

We consider an ad-hoc wireless network operating within a free market economic model. Users send data over a choice of paths, and scheduling and routing decisions are updated dynamically based on time varying channel conditions, user mobility, and current network prices charged by intermediate nodes. Each node sets its own price for relaying services, with the goal of earning revenue that exceeds its time average reception and transmission expenses. We first develop a greedy pricing strategy that maximizes social welfare while ensuring all participants make non-negative profit. We then construct a (non-greedy) policy that balances profits more evenly by optimizing a profit fairness metric. Both algorithms operate in a distributed manner and do not require knowledge of traffic rates or channel statistics. This work demonstrates that individuals can benefit from carrying wireless devices even if they are not interested in their own personal communication.

Index Terms

Revenue Maximization, Ad-Hoc Mobile, Queueing Analysis, Stochastic Optimization, Control by Pricing

I. INTRODUCTION

This paper presents a free market economic model for ad-hoc wireless networks. Multiple users desire to send packet-based traffic to their destinations, potentially using multi-hop paths. However, individual wireless nodes incur both reception and transmission costs, and hence will not agree to act as intermediate relays for this traffic unless they are adequately compensated. Thus, each node sets its own price for handling new data, and can dynamically adjust this price in reaction to time varying network conditions. Additionally, nodes dynamically choose next-hop neighbors for their data based on channel conditions and advertised prices. The goal of each user is to maximize its net utility, and the goal of individual wireless nodes is to facilitate communication while attempting to make a profit.

We design distributed pricing and control mechanisms for this system. The mechanisms yield “fruitful markets,” in the sense that the network takes maximum advantage of its multi-hop capacity while ensuring that cooperation is profitable for all participants. Specifically, we propose two different market algorithms, one that admits a greedy interpretation and one that does not.

In our first algorithm, each node charges a per-unit price that is proportional to its current level of queue backlog. Neighboring transmitters pay a *handling charge* according to this per-unit price, together with a *reception fee* that is equal to the cost incurred by receiving a new transmission. Every timeslot, individual nodes observe the current channel conditions on their outgoing links and the current prices of the corresponding neighboring nodes, and determine which data to transmit, how much to transmit, and which neighbor to send to. This choice is determined *greedily* every slot. Specifically, each node compares the past revenue earned by accepting data to the costs and service charges associated with transmitting on the current timeslot, and makes a greedy transmission decision that maximizes instantaneous profit. Note that a node might decide to remain idle on a given timeslot in order to wait for better channel conditions and/or lower prices. However, the node must eventually transmit the data, as it is obligated to remain *stable* (so that the long term output rate is equal to the long term input rate).

We show that the algorithm yields an aggregate network utility that can be pushed arbitrarily close to optimal, with a corresponding tradeoff in end-to-end average delay. Further, the algorithm ensures that

This work was presented in part at the IEEE INFOCOM conference, Anchorage, May 2007 [1].

This work is supported in part by one or both of the following: The DARPA IT-MANET program grant W911NF-07-1-0028, the National Science Foundation grant OCE 0520324.

everyone makes a non-negative profit. However, the resulting profits are not necessarily balanced evenly across members of the network. To yield a more fair profit distribution, we propose a second algorithm that seeks to maximize a general concave profit metric. This algorithm uses an interesting technique of “bang-bang pricing,” alternating between periods of allowing free service (price = 0) and periods where price is set to a pre-specified maximum value. The algorithm optimizes the target performance metric, although it relies on user cooperation and does not necessarily admit a greedy interpretation.

A. Motivations

This work treats a fundamental problem of network economics, where multiple self-interested users must coordinate to perform a routing task. Our results thus have broad applications beyond wireless scenarios. However, we focus on the ad-hoc wireless context. In this context, network infrastructure is limited and wireless users might need to relay their data through other users before entering range of a desired destination or access point. Such scenarios can occur when infrastructure never existed or was destroyed or overloaded. In such cases, privately owned wireless devices might interact with each other and with an organized team of other wireless nodes to facilitate communication. The profit earned by such interaction can encourage others to enter the market, which in turn increases network capabilities.

B. Related Work

Prior work in the area of network pricing is found in [2]-[27]. The problem of allocating flow rates to multiple users sharing a fixed capacity transmission link is considered from an economic perspective in [2]. Flow allocation and pricing in a multi-hop network is considered in [3]. Both [2] and [3] cast the problem as a static convex program, where Lagrange multipliers are interpreted as prices charged by the link to each user. It is shown that there exist prices that yield the optimal flow rates if users greedily maximize their utility minus cost. Control mechanisms that use price updates to converge to the utility optimal flow rates are considered in [4] [5]. Pricing solutions applied to static wireless downlinks are considered in [6]. More recent work in [7] [8] [30] uses back-pressure techniques for utility optimization in *stochastic* wireless networks, and relates queue backlog to prices charged to users at each network access point. Related work is considered in [9]. Worst-case throughput utility results for a wireless link with non-stationary channels are presented in [10].

In all of the works above, pricing is introduced only to obtain a fair sharing of resources over all users, so that individual profit objectives are not directly considered. Problems of pricing to maximize revenue are considered in [11] [12] [13] for static wireless downlinks, where structural properties of the resulting (non-convex) problem are examined. Work in [14] [15] [16] considers game theory approaches to related problems. Work in [17] considers admission pricing to maximize revenue in a data link with multiple traffic classes, and develops an optimal algorithm based on dynamic programming. Simplifications for large networks are considered in [18]. Market mechanisms to stimulate cooperation in *ad-hoc wireless networks* are considered in [19]-[22]. The mechanisms in [19] provide monetary credits to each node that forwards traffic, but does not consider utility optimization and does not account for heterogeneous network conditions. Work in [20] presents a simulation study of more general pricing strategies, based on pricing principles of [4]. Related work in [21] considers sub-contracting strategies for distributing a computational task over a mobile network. Analytical properties of pricing mechanisms for communication in static networks with fixed routes are considered in [22].

Our approach in this paper is quite different than the previous work, particularly that of [19]-[22], in that it provides analytical guarantees for market mechanisms, and is directly designed for *stochastic networks*. We treat the problems of social welfare and profit balancing, and ensure that all nodes make a non-negative profit by intelligently reacting to link and price information. Different from most work in this area, we consider a packet based model that fully includes queueing. Our analysis and routing strategies are inspired by the *back-pressure* concepts developed in [28] [29] and by the techniques for stochastic network optimization developed in [7] [8] [30].

C. Additional Simplifying Assumptions

While our analysis can be applied to wireless networks with general interference properties, for simplicity of exposition we consider a simplified model where each node transmits using signals that are orthogonal to those of neighboring nodes. This highlights the economic issues involved in making transmission decisions based on advertised prices and observed channel conditions, without requiring additional distributed multiple access protocols to implement these decisions. However, the multiple access problem is another important issue for wireless networks, and we briefly describe how random and scheduled access strategies can be incorporated. Specifically, suppose there are one or more owners of different network regions, and these owners schedule transmissions based on requests from nodes within their regions. We can show that the analytical results of the greedy algorithm presented in this paper are preserved if nodes pay owners a fixed fraction of their profits, so that each owner has an incentive to schedule to maximize the sum of instantaneous profit within its region. This is considered in more detail in Section V. Alternative multiple access strategies based on localized auctions within different network regions (using recent network auction results such as [23]-[27]), may also provide efficient mechanisms and suggest possible directions for future work.

In the next section we describe our network model. Sections III and IV develop the greedy and profit-balanced algorithms, respectively. Simulations for fixed-topology networks and mobile networks are provided in Section V.

II. NETWORK MODEL

Consider an ad-hoc wireless network with N nodes. The network operates in slotted time with slots $t \in \{0, 1, 2, \dots\}$. Channel conditions on each link are assumed to be constant over the duration of a timeslot, but can vary from slot to slot (due, for example, to wireless fading and/or user mobility). Specifically, let (n, b) represent the wireless link from node n to node b , and let $S_{nb}(t)$ represent the *current channel state* of the link. The value of $S_{nb}(t)$ can represent a quantized estimate of one or more physical link parameters (such as attenuation), or can represent an abstract characterization of the channel (such as “Good,” “Medium,” “Bad,” or “0”). We assume that there are a finite (but arbitrarily large) number of channel states, and that each node n knows the state of its own outgoing links at the beginning of each timeslot.

Let $\mathbf{S}(t) = (S_{nb}(t))$ represent the matrix of channel states over all network links. For simplicity of exposition, we assume throughout that channel state matrices $\mathbf{S}(t)$ are independent and identically distributed (i.i.d) over timeslots.¹ For each matrix \mathbf{S} we define channel probabilities $\pi_{\mathbf{S}} \triangleq \Pr[\mathbf{S}(t) = \mathbf{S}]$. The channel probabilities are not necessarily known to the network nodes.

A. Resource Allocation Constraints and Cost Externalities

Let $\mathbf{S}_n(t) = (S_{n1}(t), S_{n2}(t), \dots, S_{nN}(t))$ represent the vector of channel states for outgoing links of node n . We say that $S_{nb}(t) = \text{“0”}$ if node n cannot transmit to node b during slot t . In most networks of interest, nodes can only directly communicate with a small subset of current neighbors, and so each $\mathbf{S}_n(t)$ vector typically contains only a few non-zero channel states. Define $\mu_{nb}(t)$ as the transmission rate chosen by node n for the (n, b) data link during slot t (in units of bits/slot). Let $\boldsymbol{\mu}_n(t) = (\mu_{n1}(t), \mu_{n2}(t), \dots, \mu_{nN}(t))$ represent the corresponding vector of transmission rates on outgoing links of node n . The transmission rate vector for node $n \in \{1, \dots, N\}$ is chosen every timeslot in reaction to the current channel states $\mathbf{S}_n(t)$, subject to the constraint:

$$\boldsymbol{\mu}_n(t) \in \Omega_{\mathbf{S}_n(t)}^{(n)} \quad (1)$$

¹This i.i.d. assumption simplifies analysis but is not essential, and our results can be extended to general ergodic channel processes with steady state probabilities $\pi_{\mathbf{S}}$, using the T -slot Lyapunov arguments of [29][30].

where $\Omega_{S_n}^{(n)}$ represents the compact set of all transmission rate options for node n when $S_n(t) = S_n$. We assume these sets are such that if $\mu_n \in \Omega_{S_n}^{(n)}$, then we also have $\tilde{\mu}_n \in \Omega_{S_n}^{(n)}$, where $\tilde{\mu}_n$ is any vector formed from μ_n by setting one or more entries to zero. That is, it is always possible to choose to transmit nothing over a particular link, and this choice does not reduce the rate options on other links.

Note that the constraints (1) are designed for distributed implementation, as they imply that the transmission rate options available to node n are not affected by the transmission rates chosen by other nodes $m \neq n$. This assumption is valid if all nodes use orthogonal signals, so that transmission rate choices at a particular node do not influence the options of other nodes. Alternatively, this assumption holds if there is an implicit time division multiple access structure in the network, where sets of non-interfering nodes are scheduled either periodically or pseudo-randomly, and this schedule is embedded into the channel state process $S(t)$ by artificially setting link states to zero at appropriate times. Extensions to interference networks can be treated by defining constraint sets $\Omega_{S(t)}$ specifying the set of all options for the joint rate vector $(\mu_1(t), \dots, \mu_N(t))$, as discussed in more detail in [30].

Define $C_n^{tran}(\mu_n(t), S_n(t))$ as the *external transmission cost* incurred by node n due to choosing transmission rate $\mu_n(t)$ when the channel state vector is $S_n(t)$. This cost function provides a monetary measure of the personal resources (such as power) expended by node n for this transmission decision. An example transmission cost function is given by:

$$C_n^{tran}(\mu_n, S_n) = \sum_{\{b | S_{nb} > 0\}} \frac{e^{\mu_{nb}} - 1}{S_{nb}} \quad (2)$$

which corresponds to independent outgoing links, logarithmic rate-power curves $\mu_{nb} = \log(1 + S_{nb}P_{nb})$ for each link (n, b) (where S_{nb} represents an attenuation coefficient for channel (n, b)), and costs that are directly proportional to power expenditure. An example constraint set $\Omega_{S_n}^{(n)}$ is the set of all rates that can be achieved by allocating non-zero power P_{nb} to at most one outgoing link (n, b) , where $0 \leq P_{nb} \leq P_{max}$.

Define $C_{nb}^{rec}(\mu_{nb}(t))$ as the *reception cost* incurred by node b due to receiving an incoming transmission from node n over link (n, b) . This represents the external cost expended when demodulating and processing the received signal. An example reception cost function is given by:

$$C_{nb}^{rec}(\mu_{nb}) = \begin{cases} \sigma_b & \text{if } \mu_{nb} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where σ_b is a value proportional to the power expended by node b when receiving a transmission. The structure of the $C_{nb}^{rec}(\mu_{nb})$ function can also be extended to include dependence on the channel condition $S_{nb}(t)$. We assume throughout that transmission and reception costs are zero whenever the corresponding transmission rates are zero. All costs are assumed to be non-decreasing in the transmission rate vector, and are upper bounded by finite constants.

B. Network Queueing and Routing Constraints

Data might take multi-hop paths through the network, and hence each network node maintains an internal set of queues to store data according to its final destination. Any data that is bound for a particular destination node c is labeled as *commodity c data*. Let $U_n^{(c)}(t)$ represent the amount of commodity c data currently queued in network node n (in units of bits). Node n has accepted this data and hence has responsibility for either delivering this data to its destination or delivering it to another node that accepts these responsibilities. Let $\mu_{nb}^{(c)}(t)$ represent the transmission rate offered to commodity c bits over link (n, b) during slot t . Node n chooses $\mu_{nb}^{(c)}(t)$ subject to the following *routing constraints*:

$$\sum_{c=1}^N \mu_{nb}^{(c)}(t) \leq \mu_{nb}(t) \quad \text{for all links } (n, b) \text{ and all slots } t \quad (4)$$

where $\mu_{nb}(t)$ is the rate selected for link (n, b) by the resource allocation decision at node n . This model allows for *dynamic routing* of data, as both $\mu_{nb}^{(c)}(t)$ and $\mu_{nb}(t)$ can be changed from slot to slot.

Let $R_n^{(c)}(t)$ represent the amount of new exogenous commodity c data that the user at node n admits into the network during slot t . The one-step queueing dynamics for each node n and each commodity c thus satisfies:

$$U_n^{(c)}(t+1) \leq \max[U_n^{(c)}(t) - \sum_b \mu_{nb}^{(c)}(t), 0] + \sum_a \mu_{an}^{(c)}(t) + R_n^{(c)}(t) \quad (5)$$

This is expressed as an inequality because individual nodes may not have enough commodity c data to send to node n at the full offered transmission rate $\mu_{an}^{(c)}(t)$. We assume that $U_n^{(c)}(t) = 0$ for all t , as data that reaches its destination is immediately removed from the network.

C. Data Admission and Relay Pricing

We assume that each network layer node n has either zero or one user at its transport layer. Nodes with users are *source nodes*. When the user at node n admits an amount of data $R_n^{(c)}(t)$ to the network layer, it pays this node an amount $\alpha_n^{(c)}(t)$ (in units of dollars). This amount is to be determined by the dynamic pricing rule established by node n . This distinction between “the user at node n ” and “node n ” shall be convenient, even in cases when the user in fact owns node n and hence “pays itself” for acceptance of new data.

When a given node n transmits data at rate $\mu_{nb}(t)$ to another node b during slot t , it pays both a *reception fee* and a *handling charge*. The reception fee is exactly equal to the reception cost $C_{nb}^{rec}(\mu_{nb}(t))$ incurred by node b upon receiving the transmission. The handling charge is given by $\sum_c \beta_{nb}^{(c)}(t)$, where $\beta_{nb}^{(c)}(t)$ is the charge for accepting responsibility of new commodity c data, and is determined every slot by node b . The total payment from node n to node b is thus:

$$C_{nb}^{rec}(\mu_{nb}(t)) + \sum_c \beta_{nb}^{(c)}(t)$$

We shall consider handling charges of the form $\beta_{nb}^{(c)}(t) = q_b^{(c)}(t)\mu_{nb}^{(c)}(t)$, where $q_b^{(c)}(t)$ is a per-unit price for accepting commodity c data at node b . We assume that the current price $q_b^{(c)}(t)$ is set by node b and is advertised at the beginning of the timeslot, as is the reception cost function $C_{nb}^{rec}(\mu_{nb}(t))$. In this way, a transmitting node n can assess the payments required for making a transmission decision. Note that if the reception cost functions are given by (3), then each receiver node b can communicate its fixed reception fee σ_b at time 0, and every slot it needs only to advertise its current price $q_b^{(c)}(t)$ for each commodity c .

D. Time Average Profits and the Social Welfare Objective

The user at each source node n has a *utility function* $g_n^{(c)}(r)$ that represents a monetary measure of the satisfaction it receives by sending commodity c data to its destination at a long term average rate r bits/slot. Utility functions $g_n^{(c)}(r)$ are assumed to be non-negative and concave, with bounded right derivatives. We assume each user has *elastic traffic*, in the sense that it always has an infinite reservoir of data to send, and the long term send rate can be adapted to whatever rate the network allows.² In the case when the user at node n does not desire to send any data of a particular commodity c , we set $g_n^{(c)}(r) \triangleq 0$. Without loss of generality, we treat non-source nodes as if they are sources of users with utility functions that are identically zero for every commodity c .

²Inelastic traffic can be treated via techniques of [8] [30].

For each node $n \in \{1, \dots, N\}$, define the *node profit variable* $\phi_n(t)$ as follows:

$$\begin{aligned} \phi_n(t) \triangleq & \sum_c \alpha_n^{(c)}(t) + \sum_c \sum_a \beta_{an}^{(c)}(t) - \sum_c \sum_b \beta_{nb}^{(c)}(t) \\ & - C_n^{tran}(\boldsymbol{\mu}_n(t), \mathbf{S}_n(t)) - \sum_b C_{nb}^{rec}(\mu_{nb}(t)) \end{aligned} \quad (6)$$

The value $\phi_n(t)$ represents the “instantaneous profit” (total revenue minus total cost) associated with transmission decisions in the current timeslot. The total cost includes the internal payments to other nodes as well as the external transmission and reception costs. The final term in the right hand side of (6) represents the sum of all reception fees paid by node n . The external reception costs incurred by node n and the corresponding reception fees paid to node n do not appear, as these terms exactly cancel each other. Define the *expected time average profit* of node n over t slots as follows:

$$\bar{\phi}_n(t) \triangleq \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E} \{ \phi_n(\tau) \} \quad (7)$$

Likewise, for each user $n \in \{1, \dots, N\}$ we define the *expected time average user profit* $\bar{\psi}_n(t)$ as follows:

$$\bar{\psi}_n(t) \triangleq \sum_c g_n^{(c)}(\bar{r}_n^{(c)}(t)) - \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_c \mathbb{E} \{ \alpha_n^{(c)}(t) \} \quad (8)$$

where:

$$\bar{r}_n^{(c)}(t) \triangleq \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E} \{ R_n^{(c)}(\tau) \} \quad (9)$$

That is, the user profit $\bar{\psi}_n(t)$ represents the difference between the throughput utility and the time average payments associated with user n over the course of t slots.

The individual goal of each user and each node is to maximize its own time average profit. Our overall network objective is to maximize the sum of profits over all users and all nodes. However, there is an additional constraint that *all queues of the network must be stable*.³ This ensures that the long term input rate to the network is exactly the same as the long term output rate. The following simple lemma relates this sum profit objective to maximizing social welfare.

Lemma 1: (Social Welfare) Any network control and pricing algorithm that stabilizes the network yields time average profits $\bar{\psi}_n(t)$ and $\bar{\phi}_n(t)$ that satisfy:

$$\begin{aligned} \limsup_{t \rightarrow \infty} \sum_{n=1}^N [\bar{\psi}_n(t) + \bar{\phi}_n(t)] = \\ \limsup_{t \rightarrow \infty} \left[\sum_{n,c} g_n^{(c)}(\bar{r}_n^{(c)}(t)) - \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_n \mathbb{E} \{ C_n(t) \} \right] \end{aligned} \quad (10)$$

where $C_n(t)$ represents the external cost of node n :

$$C_n(t) \triangleq C_n^{tran}(\boldsymbol{\mu}_n(t), \mathbf{S}_n(t)) + \sum_b C_{nb}^{rec}(\mu_{nb}(t))$$

We call the right hand side of (10) the *social welfare* of the network, and note that it involves only external utilities and costs. The proof of the lemma is trivial, and follows by noticing that the sum of internal payments over all nodes is exactly equal to the sum of internal revenues earned from taking these payments. Therefore, the internal monetary costs and revenues cancel each other out in the sum profit

³We say a queue $U(t)$ is stable if $\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E} \{ U(\tau) \} < \infty$. This type of stability is usually referred to as *strong stability*.

metric. The lemma implies that the objective of designing a network control and pricing algorithm to stabilize the network while maximizing sum profit is equivalent to the objective of designing a network control algorithm to maximize the social welfare metric, *without regard to network prices*.

The following theorem establishes that any achievable social welfare value (and hence any achievable sum profit value) can be achieved arbitrarily closely via a stationary randomized policy that bases decisions only on the current channel state, and that sets all monetary charges $\alpha_n^{(c)}(t)$ and $\beta_n^{(c)}(t)$ to zero.

Theorem 1: Suppose there exists a control strategy that stabilizes the network and yields a positive lim sup social welfare value g^* :

$$g^* \triangleq \limsup_{t \rightarrow \infty} \left[\sum_{n,c} g_n^{(c)}(\bar{r}_n^{(c)}(t)) - \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_n \mathbb{E} \{C_n(t)\} \right]$$

Then for any ρ such that $0 < \rho < 1$, there exists a stationary randomized control algorithm that stabilizes the network, sets all $\alpha_n^{(c)}(t), \beta_n^{(c)}(t)$ to zero, sets all admissions $R_n^{(c)}(t)$ to particular constant values $R_n^{(c)*}$ for all time, and that chooses transmission rates $\mu_n^*(t)$ according to a stationary and random function of the observed channel state matrix $\mathbf{S}(t)$. Further, this stationary randomized policy yields for all t and all (n, c) :

$$R_n^{(c)*} + \sum_a \mathbb{E} \{ \mu_{an}^{(c)*}(t) \} \leq \sum_b \mathbb{E} \{ \mu_{nb}^{(c)*}(t) \} \quad (11)$$

and yields the following social welfare result for all slots t :

$$\sum_{n,c} g_n^{(c)}(R_n^{(c)*}) - \sum_n \mathbb{E} \{C_n(t)\} \geq \rho g^* \quad (12)$$

Proof: The proof is similar to the necessary conditions for network stability and minimum average energy expenditure proven in [29] and [31], and is omitted for brevity. \square

The probabilities and transmission rate modes required of the stationary policy in Theorem 1 could in principle be computed by an offline algorithm with centralized knowledge of all channel probabilities, cost functions, and user utilities. However, the resulting algorithm might cause some nodes to receive negative profit, and hence these nodes would have no incentive to continue participating. The design of an online control algorithm that maximizes social welfare in this context and ensures all users and nodes receive non-negative profit is an open question. We resolve this question in the next section by a simple and direct online algorithm that makes use of *backpressure* [28] [29]. The algorithm has the additional desirable feature that individual control actions can be interpreted as greedily maximizing instantaneous profit.

E. Discussion of Alternative Approaches

It is possible to use a modified Lagrange multiplier argument, similar to [3], to prove existence of fixed flows and prices that yield the desired non-negative profit result in a static network with no channel variation and with more assumptions imposed on the structure of the cost functions. However, the resulting multipliers (prices) are not known a-priori, and it would require an extensive offline computation to estimate them to within an adequate degree of accuracy. Online techniques related to dual subgradient algorithms, as in [29] [8] [9], can also be considered when prices are suitably defined for this free market context. However, we can show that these algorithms *do not* always lead to non-negative profits, particularly when there are centrally located nodes with varying channels and non-negligible costs.

III. A GREEDY PRICING STRATEGY

The following algorithm makes distributed and greedy decisions at each node using local link conditions and prices of neighboring nodes. It uses a positive constant V that determines a tradeoff in end-to-end

network delay. We shall also require the following finite bounds on the maximum transmission rate into and out of a given node n :

$$\begin{aligned}\mu_n^{max,out} &\triangleq \sup_{[S, \mu_n \in \Omega_{S_n}^{(n)}]} \sum_b \mu_{nb} \\ \mu_n^{max,in} &\triangleq \sup_{[S, \{\mu_i \in \Omega_{S_i}^{(i)}\}_{i=1}^N]} \sum_a \mu_{an}\end{aligned}$$

Stochastic Greedy Pricing Algorithm (SGP):

Pricing: Every timeslot t , each node n sets the per-unit price $q_n^{(c)}(t)$ for handling new commodity c data as follows:

$$q_n^{(c)}(t) = U_n^{(c)}(t)/V$$

The corresponding charge for accepting $R_n^{(c)}(t)$ units of exogenous data is given by:

$$\alpha_n^{(c)}(t) = R_n^{(c)}(t)q_n^{(c)}(t)$$

The charge for accepting endogenous commodity c data at rate $\mu_{an}^{(c)}(t)$ is given by:

$$\beta_{an}^{(c)}(t) = \mu_{an}^{(c)}(t)q_n^{(c)}(t)$$

Admission Control: Every timeslot t , each user n observes the current prices $q_n^{(c)}(t)$ in its source node (for all commodities c such that the utility function $g_n^{(c)}(r)$ is not identically zero), and chooses $R_n^{(c)}(t) = r_n^{(c)}$, where the $r_n^{(c)}$ values solve:

$$\begin{aligned}\text{Maximize: } & \sum_c g_n^{(c)}(r_n^{(c)}) - \sum_c r_n^{(c)} q_n^{(c)}(t) \\ \text{Subject to: } & \sum_c r_n^{(c)} \leq R_n^{max} \\ & 0 \leq r_n^{(c)} \text{ for all } c\end{aligned}\tag{13}$$

where R_n^{max} is a constant such that $R_n^{max} \geq \mu_n^{max,out}$.

Resource Allocation: Every timeslot t , each node n observes the current prices $q_b^{(c)}(t)$ advertised by neighboring nodes b . It then computes the *differential price* $W_{nb}^{(c)}(t)$ as follows:

$$W_{nb}^{(c)}(t) \triangleq q_n^{(c)}(t) - q_b^{(c)}(t) - \delta_{max}/V\tag{14}$$

where $\delta_{max} \triangleq \max_n \{\mu_n^{max,out}, \mu_n^{max,in} + R_n^{max}\}$, and represents the largest change in any queue backlog during a slot. The network parameter δ_{max} is assumed to be known by all nodes at the beginning of operation. The *optimal* differential price and the corresponding optimal commodity is computed:

$$\begin{aligned}W_{nb}^*(t) &\triangleq \max_c W_{nb}^{(c)}(t) \\ c_{nb}^*(t) &\triangleq \arg \max_c W_{nb}^{(c)}(t)\end{aligned}\tag{15}$$

The current channel states $S_n(t)$ of all outgoing links are observed, and the transmission rate vector $\mu_n(t)$ is allocated as the solution of the following optimization problem:

$$\begin{aligned}\text{Maximize: } & \sum_b \mu_{nb} W_{nb}^*(t) - \sum_b C_{nb}^{rec}(\mu_{nb}) \\ & - C_n^{tran}(\mu_n, S_n(t)) \\ \text{Subject to: } & \mu_n \in \Omega_{S_n(t)}^{(n)}\end{aligned}\tag{16}$$

Routing/Scheduling: Whenever $W_{nb}^*(t) > 0$, data of commodity $c_{nb}^*(t)$ is transmitted over link (n, b) at a rate $\mu_{nb}(t)$, where $\mu_{nb}(t)$ is determined by the resource allocation algorithm above.

A. Greedy Interpretation of SGP

The admission control strategy can be viewed as a greedy optimization every timeslot, where node n maximizes its utility associated with admitting new traffic to its source node (as measured by the $g_n^{(c)}(r)$ functions) minus the total cost of admitting this traffic, subject to a constraint R_n^{max} on the total sum of admitted data. Note that in the special case when user n has only a single active commodity (that is, it has a single non-zero $g_n^{(c)}(r)$ function), then the algorithm reduces to setting $R_n^{(c)}(t) = r$, where r maximizes $g_n^{(c)}(r) - r q_n^{(c)}(t)$ subject to $0 \leq r \leq R_n^{max}$.

The resource allocation computes the differential prices $W_{nb}^{(c)}(t)$. Note from (14) that:

$$W_{nb}^{(c)}(t) \mu_{nb}^{(c)}(t) = [q_n^{(c)}(t) - q_b^{(c)}(t)] \mu_{nb}^{(c)}(t) - \mu_{nb}^{(c)}(t) \delta_{max}/V$$

The first term on the right hand side of the above equality represents the difference between the charge required for transmitting commodity c traffic to node b and the revenue it would earn by accepting this same amount of traffic from other nodes, charging these nodes the current price $q_n^{(c)}(t)$. From a greedy perspective, it makes sense to transmit commodity c data to node b only when the price differential $q_n^{(c)}(t) - q_b^{(c)}(t)$ is positive. The value $\mu_{nb}^{(c)}(t) \delta_{max}/V$ can be viewed as a “fudge factor” that decreases the price differential to account for the fact that node n may not receive new data at its currently advertised price. The value $W_{nb}^{(c)}(t) \mu_{nb}^{(c)}(t)$ can thus be viewed as node n ’s estimate of its instantaneous profit associated with relaying $\mu_{nb}^{(c)}(t)$ units of commodity c data (not including transmission costs or reception fees). Hence, commodity $c_{nb}^*(t)$ defined in (15) is the most valuable commodity to transfer over link (n, b) . The resource allocation (16) can thus be viewed as a greedy attempt by node n to allocate resources to maximize its total instantaneous profit.

Note also that the SGP algorithm transmits commodity c data from node n only if $W_{nb}^{(c)}(t) > 0$ for some receiver node b . It follows from (14) that such a transmission can only take place if $U_n^{(c)}(t) > \delta_{max}$ (recall that SGP uses price $q_n^{(c)}(t) = U_n^{(c)}(t)/V$). Therefore, there is always enough data available to fill the offered transmission rates.

B. Algorithm Performance

Assume all queues of the network are initially empty, and that the SGP algorithm is implemented with a fixed control parameter $V > 0$. Assume all utility functions $g_n^{(c)}(r)$ have finite right derivatives at $r = 0$. Because utilities are concave, it follows that right derivatives exist and are non-increasing over the interval $r \geq 0$. Define η as the maximum right derivative of any utility function. Note that $0 \leq \eta < \infty$. Additionally assume that utility is zero when rate is zero, i.e., $g_n^{(c)}(0) = 0$ for all n and c .

Theorem 2: (SGP Performance) For arbitrary $S(t)$ processes and for any fixed parameter $V > 0$, the SGP algorithm ensures:

(a) $U_n^{(c)}(t) \leq V\eta + \delta_{max}$ for all slots t and all (n, c) .

(b) All nodes and users receive non-negative profit. Specifically, for all slots t and nodes $n \in \{1, \dots, N\}$, we have:

$$\frac{1}{t} \sum_{\tau=0}^{t-1} \phi_n(\tau) \geq 0 \quad (17)$$

Likewise, for all users $n \in \{1, \dots, N\}$, all commodities c , and all slots t , we have:

$$g_n^{(c)} \left(\frac{1}{t} \sum_{\tau=0}^{t-1} R_n^{(c)}(\tau) \right) - \frac{1}{t} \sum_{\tau=0}^{t-1} \alpha_n^{(c)}(\tau) \geq 0 \quad (18)$$

(c) If channel state matrices $\mathbf{S}(t)$ are i.i.d. over timeslots, then the achieved social welfare of the algorithm satisfies:

$$\liminf_{t \rightarrow \infty} \left[\sum_{n,c} g_n^{(c)}(\bar{r}_n^{(c)}(t)) - \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_n \mathbb{E} \{C_n(t)\} \right] \geq g^* - O(1/V)$$

where g^* is the corresponding social welfare value achieved by any other stabilizing control algorithm.⁴

The parameter V thus determines an explicit tradeoff between welfare utility and queue congestion (and hence, by Little's Theorem [32], end-to-end average delay). It is interesting to note that the non-negative profit result of part (b) holds deterministically on every timeslot t and for any node n that is implementing the SGP algorithm, regardless of whether or not the other nodes are implementing SGP.

Proof: (Theorem 2 part (a)) Fix any (n, c) pair, and consider the admission variable $R_n^{(c)}(t)$ chosen by the SGP algorithm according to (13) at a particular time t . The right derivative of $g_n^{(c)}(r)$ evaluated at any point $r \geq 0$ is less than or equal to η . Hence, if $q_n^{(c)}(t) > \eta$, then $g_n^{(c)}(r) - r q_n^{(c)}(t) \leq 0$ for all $r \geq 0$, with equality holding only at $r = 0$ (recall that $g_n^{(c)}(0) = 0$). It follows that if $q_n^{(c)}(t) > \eta$, then $R_n^{(c)}(t) = 0$ (otherwise, the solution to (13) could be improved by setting $R_n^{(c)}(t) = 0$). Noting that $q_n^{(c)}(t) = U_n^{(c)}(t)/V$, we have established the following important property (Property P1) of SGP:

(P1) For any $(n, c), t$, if $U_n^{(c)}(t) > V\eta$, then $R_n^{(c)}(t) = 0$.

Now suppose that for a particular timeslot t , we have $U_n^{(c)}(t) \leq V\eta + \delta_{max}$ for all (n, c) (this certainly holds for $t = 0$, as all queues are initially empty). We prove that the same holds for time $t + 1$.

Consider any particular (n, c) . If $U_n^{(c)}(t + 1) \leq U_n^{(c)}(t)$, then clearly $U_n^{(c)}(t + 1) \leq V\eta + \delta_{max}$. Else, queue (n, c) must have received new commodity c arrivals during slot t (either endogenous, exogenous, or both). If it received a positive amount of exogenous commodity c arrivals from the source user, then $R_n^{(c)}(t) > 0$ (refer to the one-step queueing dynamics (5)). By Property P1, this implies that $U_n^{(c)}(t) \leq V\eta$. As δ_{max} represents the largest change in queue backlog during any single timeslot, it follows that $U_n^{(c)}(t + 1) \leq V\eta + \delta_{max}$.

Finally, in the case that this queue did not receive any exogenous arrivals but did receive a positive amount of endogenous data transmitted from other nodes, then by the SGP routing policy we know that $W_{an}^{(c)}(t) > 0$ for at least one other node a (where $a \neq n$). It follows from (14) and the price definitions $q_n^{(c)}(t) = U_n^{(c)}(t)/V$ that:

$$U_a^{(c)}(t) - U_n^{(c)}(t) - \delta_{max} > 0 \quad (19)$$

Therefore:

$$\begin{aligned} U_n^{(c)}(t + 1) &\leq U_n^{(c)}(t) + \delta_{max} \\ &< U_a^{(c)}(t) - \delta_{max} + \delta_{max} \end{aligned} \quad (20)$$

$$\leq V\eta + \delta_{max} \quad (21)$$

where (20) follows from (19), and (21) follows because all queues are bounded by $V\eta + \delta_{max}$ on slot t .

Hence, in all cases we have $U_n^{(c)}(t + 1) \leq V\eta + \delta_{max}$. This holds for all queues (n, c) , and by induction it holds for all timeslots $t \in \{0, 1, 2, \dots\}$, proving the result. \square

Proof: (Theorem 2 part (b)) For each user n , the SGP admission decisions $R_n^{(c)}(t)$ are chosen to optimally solve (13). Hence, for any (n, c) and any slot τ we have:

$$g_n^{(c)}(R_n^{(c)}(\tau)) - q_n^{(c)}(\tau) R_n^{(c)}(\tau) \geq 0 \quad (22)$$

Indeed, the left hand side being negative would create a contradiction, as then the solution to (13) could be strictly improved by changing $R_n^{(c)}(\tau)$ to 0. Taking a time average of (22) over $\tau \in \{0, \dots, t - 1\}$ and using concavity of the utility function together with the fact that $\alpha_n^{(c)}(\tau) = q_n^{(c)}(\tau) R_n^{(c)}(\tau)$ yields (18).

⁴A bound on the $O(1/V)$ term in the theorem can be computed explicitly, but we omit this computation for brevity.

To prove (17), assume that a given node n implements SGP. For each slot τ , define $h_n(\tau)$ as follows:

$$h_n(\tau) \triangleq \sum_b \sum_c \mu_{nb}^{(c)}(\tau) W_{nb}^{(c)}(\tau) - \sum_b C_{nb}^{rec}(\mu_{nb}(\tau)) - C_n^{tran}(\mu_n(\tau), \mathbf{S}_n(\tau)) \quad (23)$$

For any link (n, b) , SGP transmits only a single commodity $c_{nb}^*(\tau)$, and this commodity receives the full transmission rate $\mu_{nb}(\tau)$. It follows that $\mu_{nb}(\tau) W_{nb}^*(\tau) = \sum_c \mu_{nb}^{(c)}(\tau) W_{nb}^{(c)}(\tau)$. Thus, $h_n(\tau)$ is the same as the maximization metric (16) used in the resource allocation algorithm of SGP. Further note that $\mu_n(\tau) = \mathbf{0}$ is always an option in the resource allocation optimization (16), and hence this optimization metric is always non-negative. That is, $h_n(\tau) \geq 0$ for all τ . Using simple arithmetic together with the definitions of $h_n(\tau)$ in (23) and $\phi_n(\tau)$ in (6), we have for all τ :

$$\begin{aligned} \phi_n(\tau) &= h_n(\tau) + \sum_c \alpha_n^{(c)}(\tau) + \sum_c \sum_a \beta_{an}^{(c)}(\tau) - \sum_c \sum_b \beta_{nb}^{(c)}(\tau) \\ &\quad - \sum_{b,c} \mu_{nb}^{(c)}(\tau) [q_n^{(c)}(\tau) - q_b^{(c)}(\tau) - \delta_{max}/V] \\ &= h_n(\tau) + \sum_c \alpha_n^{(c)}(\tau) + \sum_c \sum_a \beta_{an}^{(c)}(\tau) \\ &\quad - \sum_{b,c} \mu_{nb}^{(c)}(\tau) [q_n^{(c)}(\tau) - \delta_{max}/V] \end{aligned}$$

where the final equality holds because, by definition, $\beta_{nb}^{(c)}(\tau) = \mu_{nb}^{(c)}(\tau) q_b^{(c)}(\tau)$. Because $h_n(\tau) \geq 0$, we have:

$$\begin{aligned} \sum_{\tau=0}^{t-1} \phi_n(\tau) &\geq \sum_{\tau=0}^{t-1} \left[\sum_c \alpha_n^{(c)}(\tau) + \sum_c \sum_a \beta_{an}^{(c)}(\tau) \right] \\ &\quad - \sum_{\tau=0}^{t-1} \sum_{b,c} \mu_{nb}^{(c)}(\tau) [q_n^{(c)}(\tau) - \delta_{max}/V] \end{aligned} \quad (24)$$

It suffices to prove that the right hand side of the above inequality is non-negative. That is, we desire to prove:

$$Revenue_n(t) \geq \sum_{\tau=0}^{t-1} \sum_{b,c} \mu_{nb}^{(c)}(\tau) [q_n^{(c)}(\tau) - \delta_{max}/V] \quad (25)$$

where $Revenue_n(t)$ represents the first term on the right hand side of (24), and is equal to the total revenue earned by node n from handling charges paid to it during the course of the first t slots.

To show this, observe that all sample paths and queueing values are preserved if the actual data chosen to be transmitted from each queue $U_n^{(c)}(\tau)$ takes place according to the *Last In First Out* (LIFO) strategy. Using this interpretation, we note that every bit of data that arrives to node n is charged a particular price by node n . Thus, for a particular slot τ , the data associated with the $\mu_{nb}^{(c)}(\tau)$ transmissions *out* of node n is composed of bits that may have arrived at different times and may have been charged different prices. However, under the LIFO transmission rule, all of this data was transmitted *into* node n during slots when the queue backlog was greater than or equal to $U_n^{(c)}(\tau) - \delta_{max}$. This is because $\mu_{nb}^{(c)}(\tau) \leq \delta_{max}$, and so the LIFO transmission of $\mu_{nb}^{(c)}(\tau)$ data leaves at least $U_n^{(c)}(\tau) - \delta_{max}$ bits of data behind, all of which must have been there when the transmitted data arrived. It follows that the price charged to each bit of this transmitted data when it arrived to node n was *at least* $(U_n^{(c)}(\tau) - \delta_{max})/V$. This is because node n uses the SGP pricing rule that sets current price equal to current backlog divided by V , and holds regardless

of the pricing and decision strategies of the other nodes. Therefore, the total revenue earned by accepting this data is at least as large as the right hand side of (25). Thus, (25) holds, and the result follows. \square

Part (c) of Theorem 2 is derived via the following Lyapunov drift lemma from [8][30]. Let $\mathbf{U}(t) = (U_n^{(c)}(t))$ represent the matrix of queue backlogs, and let $L(\mathbf{U})$ be a non-negative function of the network queue state, called a *Lyapunov function*. Formally define the *Lyapunov drift* as follows:

$$\Delta(\mathbf{U}(t)) \triangleq \mathbb{E} \{L(\mathbf{U}(t+1)) - L(\mathbf{U}(t)) \mid \mathbf{U}(t)\}$$

Let $f(t)$ represent some real valued stochastic reward process related to the system, and assume $f(t) \leq f_{max}$ for all t , for some finite upper bound f_{max} .

Lemma 2: (Lyapunov drift [8][30]) If there exist constants $B > 0$, $\epsilon > 0$, $V > 0$ such that for all slots t and all queue states $\mathbf{U}(t)$, the Lyapunov drift satisfies:

$$\Delta(\mathbf{U}(t)) - V \mathbb{E} \{f(t) \mid \mathbf{U}(t)\} \leq B - \epsilon \sum_{n,c} U_n^{(c)}(t) - V f^*$$

for some target utility value f^* , then:

$$\begin{aligned} \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_{n,c} \mathbb{E} \{U_n^{(c)}(\tau)\} &\leq \frac{B + V(f_{max} - f^*)}{\epsilon} \\ \liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E} \{f(\tau)\} &\geq f^* - B/V \quad \square \end{aligned}$$

Proof: (Theorem 2 part (c)) Define $L(\mathbf{U}) \triangleq \frac{1}{2} \sum_{n,c} (U_n^{(c)})^2$. The queueing dynamics (5) can be used to bound the Lyapunov drift $\Delta(\mathbf{U}(t))$ according to a standard computation [29][30]:

$$\begin{aligned} \Delta(\mathbf{U}(t)) - V \sum_n \mathbb{E} \left\{ \sum_c g_n^{(c)}(R_n^{(c)}(t)) - C_n(t) \mid \mathbf{U}(t) \right\} &\leq \\ B - V \sum_n \mathbb{E} \left\{ \sum_c g_n^{(c)}(R_n^{(c)}(t)) - C_n(t) \mid \mathbf{U}(t) \right\} & \\ - \sum_{n,c} U_n^{(c)}(t) \mathbb{E} \left\{ \sum_b \mu_{nb}^{(c)}(t) - \sum_a \mu_{an}^{(c)}(t) - R_n^{(c)}(t) \mid \mathbf{U}(t) \right\} & \end{aligned}$$

where B is a constant that depends on N and δ_{max} .⁵ The key observation is that, given $\mathbf{U}(t)$ and given the pricing rule $q_n^{(c)}(t) = U_n^{(c)}(t)/V$, the SGP algorithm *comes within an additive constant* $\tilde{B} \triangleq \delta_{max} \sum_n \mu_n^{max,out}$ of minimizing the right hand side of the above inequality over all possible control decisions for $\{R_n^{(c)}(t)\}$, $\{\mu_n(t)\}$, and $\{\mu_{nb}^{(c)}(t)\}$. The additive constant is because we use the modified differential backlog in (14) (with a δ_{max}/V “fudge factor”), rather than the pure differential backlog. The detailed demonstration of this is similar to related demonstrations in [8][30], and is omitted for brevity. However, this immediately implies that plugging the *alternative* stationary randomized decisions $\{R_n^{(c)*}\}$, $\{\mu_n^*(t)\}$, and $\{\mu_{nb}^{(c)*}(t)\}$ from Theorem 1 into the right hand side and adding \tilde{B} preserves the inequality. The stationary randomized algorithm makes decisions independent of $\mathbf{U}(t)$, and yields a significant simplification. Specifically, plugging (11) and (12) directly into the right hand side of the above inequality and adding \tilde{B} yields:

$$\begin{aligned} \Delta(\mathbf{U}(t)) - V \sum_n \mathbb{E} \left\{ \sum_c g_n^{(c)}(R_n^{(c)}(t)) - C_n(t) \mid \mathbf{U}(t) \right\} &\leq \\ (B + \tilde{B}) - V \rho g^* & \end{aligned}$$

⁵Specifically, it can be shown that $B \leq \frac{1}{2} \sum_n (\mu_n^{max,out})^2 + \frac{1}{2} \sum_n (\mu_n^{max,in} + R_n^{max})^2$, as in [29].

where g^* is the social welfare of any particular stabilizing strategy, and ρ is any value such that $0 < \rho < 1$ (from Theorem 1). Using Lemma 2 yields:

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_n \mathbb{E} \left\{ \sum_c g_n^{(c)}(R_n^{(c)}(\tau)) - C_n(\tau) \right\} \geq \rho g^* - (B + \tilde{B})/V$$

Taking a limit as $\rho \rightarrow 1$ and using concavity of the utility functions $g_n^{(c)}(r)$ proves part (c) of Theorem 2. \square

IV. PRICING FOR BALANCED PROFITS

While the SGP algorithm makes greedy decisions and ensures non-negative profit for all participants, the profits might not be distributed evenly. To provide more balanced profits, we define *profit metrics* according to general non-decreasing concave functions $\Phi_n(\phi_n)$ and $\Psi_n(\psi_n)$ for all nodes and users. Suppose that $\bar{\phi}_n$ and $\bar{\psi}_n$ respectively represent the time average profit of each node and user $n \in \{1, \dots, N\}$. The goal is to design a control and pricing algorithm that stabilizes the network, ensures non-negative profits, and that optimizes:

$$\sum_{n=1}^N [\Phi_n(\bar{\phi}_n) + \Psi_n(\bar{\psi}_n)] \quad (26)$$

We impose an additional *bounded price assumption* on the problem: all per-unit prices are bounded by a maximum price Q_{max} . To solve this problem, we use our stochastic network optimization framework from [30] [31] [8]. Let $\gamma_n(t)$ and $\nu_n(t)$ represent *auxiliary variables*, and consider the equivalent problem:

$$\begin{aligned} \text{Maximize:} \quad & \sum_{n=1}^N [\Phi_n(\bar{\gamma}_n) + \Psi_n(\bar{\nu}_n)] \\ \text{Subject to:} \quad & 1) \quad \bar{\phi}_n \geq \bar{\gamma}_n, \quad \bar{\psi}_n \geq \bar{\nu}_n \quad \text{for all } n \\ & 2) \quad \text{Network Stability} \end{aligned}$$

That this new problem is equivalent to the problem of maximizing (26) subject to network stability follows immediately from the fact that the $\Phi_n(\phi)$ and $\Psi_n(\psi)$ functions are non-decreasing. The intuition behind why it is important to modify the problem by introducing auxiliary variables in this way is the following: The auxiliary variables $\gamma_n(t)$ and $\nu_n(t)$ are chosen in software and can be set to any desired value at any time, regardless of the network state. Thus, it is easier to use such auxiliary variables in the maximization of the concave (and hence potentially non-linear) functions $\Phi_n()$ and $\Psi_n()$. In contrast, the actual user and node profit variables $\phi_n(t)$ and $\psi_n(t)$ cannot be chosen as desired every slot, as these instantaneous profits depend highly on the current network state. Thus, it would be difficult (or impossible) to use $\phi_n(t)$ and $\psi_n(t)$ directly in an on-line nonlinear optimization every slot, but it is possible to ensure that their time averages conform to the *linear* time average inequality constraints $\bar{\phi}_n \geq \bar{\gamma}_n$ and $\bar{\psi}_n \geq \bar{\nu}_n$ (further details on this important stochastic optimization technique are discussed in [30]). This issue did not arise in our previous analysis of SGP, as the variables that appeared in the nonlinear optimization metrics were admission decisions $R_n^{(c)}(t)$, which could be chosen as desired each slot because of the elastic traffic assumption that users always have data to send.⁶

Define $\phi_n^{pos}(t)$ and $\phi_n^{neg}(t)$ respectively as the sum of positive terms in the $\phi_n(t)$ function (6) and the absolute value of the sum of the negative terms. Thus, $\phi_n(t) = \phi_n^{pos}(t) - \phi_n^{neg}(t)$. The inequality constraints

⁶Indeed, without assuming users have infinite backlog for admission decisions (as, for example, when new data is not always available but arrives randomly to the transport layer at each source), the SGP algorithm would need to be modified by introducing auxiliary variables or “flow state variables,” a technique developed in a related flow control context (without profit metrics) in [8] [30].

in the above problem are treated via stabilization of a new set of *virtual queues* $X_n(t)$ and $Y_n(t)$, defined according to the following dynamics:

$$\begin{aligned} X_n(t+1) &= \max[X_n(t) - \phi_n^{pos}(t), 0] + \gamma_n(t) + \phi_n^{neg}(t) \\ Y_n(t+1) &= \max[Y_n(t) - \sum_c g_n^{(c)}(R_n^{(c)}(t)), 0] \\ &\quad + \nu_n(t) + \sum_c \alpha_n^{(c)}(t) \end{aligned} \quad (27)$$

Stabilizing the $X_n(t)$ and $Y_n(t)$ queues implies that the time average of the queue input variables is less than or equal to the time average of the queue service variables. Specifically, stabilizing $X_n(t)$ ensures $\bar{\gamma}_n + \bar{\phi}_n^{neg} \leq \bar{\phi}_n^{pos}$, which is equivalent to the constraint $\bar{\phi}_n \leq \bar{\gamma}_n$. Likewise, stabilizing $Y_n(t)$ ensures that $\bar{\nu}_n \leq \sum_c [\bar{g}_n^{(c)}(\bar{r}_n^{(c)}) - \bar{\alpha}_n^{(c)}]$. We next define the pricing variables:

$$\alpha_n^{(c)}(t) = p_n^{(c)}(t) R_n^{(c)}(t), \quad \beta_{nb}^{(c)}(t) = q_{nb}^{(c)}(t) \mu_{nb}^{(c)}(t)$$

where $p_n^{(c)}(t)$ is the per-unit price for exogenous commodity c arrivals at node n , and $q_{nb}^{(c)}(t)$ is the per-unit price for commodity c data transmitted over link (n, b) . Recall that in SGP we found it was sufficient to use a single price $q_n^{(c)}(t)$ that is charged to any user or incoming link that sends new commodity c data to node n on slot t (so that $p_n^{(c)}(t) = q_{an}^{(c)}(t) = q_n^{(c)}(t)$ for each link (a, n)). However, here we find it is important to allow the possibility of differentiated pricing for each link and user sending data to node n .

Let $\mathbf{Z}(t) \triangleq (\mathbf{U}(t); \mathbf{X}(t); \mathbf{Y}(t))$ represent the combined queue state. Define the Lyapunov function $L(\mathbf{Z}(t))$ as follows:

$$L(\mathbf{Z}(t)) = \frac{1}{2} \sum_{n,c} (U_n^{(c)}(t))^2 + \frac{1}{2} \sum_n X_n(t)^2 + \frac{1}{2} \sum_n Y_n(t)^2$$

The following bound involving Lyapunov drift $\Delta(\mathbf{Z}(t))$ can be computed via the queueing dynamics (5) and (27):

$$\begin{aligned} \Delta(\mathbf{Z}(t)) - V \sum_n \mathbb{E} \{ \Phi_n(\gamma_n(t)) + \Psi_n(\nu_n(t)) \mid \mathbf{Z}(t) \} \leq \\ D - V \sum_n \mathbb{E} \{ \Phi_n(\gamma_n(t)) + \Psi_n(\nu_n(t)) \mid \mathbf{Z}(t) \} \\ - \sum_{n,c} U_n^{(c)}(t) \mathbb{E} \left\{ \sum_b \mu_{nb}^{(c)}(t) - \sum_a \mu_{an}^{(c)}(t) - R_n^{(c)}(t) \mid \mathbf{Z}(t) \right\} \\ - \sum_n X_n(t) \mathbb{E} \{ \phi_n(t) - \gamma_n(t) \mid \mathbf{Z}(t) \} \\ - \sum_n Y_n(t) \mathbb{E} \left\{ \sum_c g_n^{(c)}(R_n^{(c)}(t)) - \sum_c \alpha_n^{(c)}(t) - \nu_n(t) \mid \mathbf{Z}(t) \right\} \end{aligned}$$

where D is a constant. The following algorithm is obtained by making control and pricing decisions that minimize the right hand side of the above drift bound on every slot t .

Bang-Bang Pricing Algorithm for Stochastic Networks:

Pricing: Every slot t , each node n observes its virtual queues $X_n(t)$, $Y_n(t)$ and chooses $p_n^{(c)}(t)$ as follows:

$$p_n^{(c)}(t) = \begin{cases} Q_{max} & \text{if } Y_n(t) < X_n(t) \\ 0 & \text{otherwise} \end{cases}$$

Each node n also observes the virtual queues $X_a(t)$ of its neighbors, and chooses $q_{an}^{(c)}(t)$ as follows:

$$q_{an}^{(c)}(t) = \begin{cases} Q_{max} & \text{if } X_a(t) < X_n(t) \\ 0 & \text{otherwise} \end{cases}$$

Admission Control: Every slot t , each node n chooses $R_n^{(c)}(t)$ (for each commodity c) as the maximum of:

$$Y_n(t) g_n^{(c)}(r) - r [U_n^{(c)}(t) - p_n^{(c)}(t) (X_n(t) - Y_n(t))]$$

over the interval $0 \leq r \leq R_n^{max}$.

For each n , the auxiliary variable $\gamma_n(t)$ is chosen as the maximum of $V\Phi_n(\gamma) - X_n(t)\gamma$ subject to $0 \leq \gamma \leq Q_{max}\delta_{max}$. Likewise, the auxiliary variable $\nu_n(t)$ is chosen as the maximum of $V\Psi_n(\nu) - Y_n(t)\nu$ subject to $0 \leq \nu \leq \eta R_n^{max}$.

Resource Allocation: Each node n observes the channel states and queue backlogs of its neighbors, and computes:

$$\Theta_{nb}^{(c)}(t) \triangleq U_n^{(c)}(t) - U_b^{(c)}(t) - q_{nb}^{(c)}(t)(X_n(t) - X_b(t))$$

The optimal weight and commodity is then chosen as follows:

$$\Theta_{nb}^*(t) \triangleq \max_c \Theta_{nb}^{(c)}(t), \quad c_{nb}^*(t) \triangleq \arg \max_c \Theta_{nb}^{(c)}(t)$$

The transmission vector $\mu_n(t)$ is allocated as the solution to:

$$\begin{aligned} \text{Maximize:} \quad & \sum_b \mu_{nb} \Theta_{nb}^*(t) - X_n(t) C_n^{tran}(\mu_n, \mathbf{S}_n(t)) \\ & - X_n(t) \sum_b C_{nb}^{rec}(\mu_{nb}) \\ \text{Subject to:} \quad & \mu_n \in \Omega_{\mathbf{S}_n(t)}^{(n)} \end{aligned}$$

Routing/Scheduling: The $\mu_{nb}^{(c)}(t)$ rates are selected as:

$$\mu_{nb}^{(c)}(t) = \begin{cases} \mu_{nb}(t) & \text{if } \Theta_{nb}^*(t) > 0 \text{ and } c = c_{nb}^*(t) \\ 0 & \text{otherwise} \end{cases}$$

The actual queues $\mathbf{U}(t)$ are then updated according to (5), and the virtual queues $\mathbf{X}(t), \mathbf{Y}(t)$ are updated according to (27).

Unlike SGP, the pricing here does not depend on the commodity c , but is potentially link dependent.

Theorem 3: (Bang-Bang Pricing Performance) The Bang-Bang pricing algorithm stabilizes all actual and virtual queues of the system, ensures all participants make non-negative time average profit, and yields:

$$\begin{aligned} \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_{n,c} \mathbb{E} \{U_n^{(c)}(\tau)\} & \leq O(V) \\ \liminf_{t \rightarrow \infty} \sum_{n=1}^N [\Phi_n(\bar{\phi}_n(t)) + \Psi_n(\bar{\psi}_n(t))] & \geq \Phi^* - O(1/V) \end{aligned}$$

where $\bar{\phi}_n(t)$ and $\bar{\psi}_n(t)$ are defined in (7) and (8), and where Φ^* is the lim sup of the achieved profit metric (26) under any other stabilizing control algorithm.

Proof: The result uses the Lyapunov drift lemma (Lemma 2) together with the fact that (i) The given control decision variables minimize the right hand side of the drift bound, and (ii) There exists a stationary randomized control algorithm that stabilizes the system and achieves a profit metric of at least $\rho\Phi^*$ (for any ρ such that $0 < \rho < 1$). A complete derivation is omitted for brevity. \square

V. SIMULATION

A. A Fixed-Topology Network with Access Points

Consider the network of Fig. 1, where there are seven wireless nodes, four of which have sources that desire to send data to any of the three wireline access points. This is a single commodity problem, as the three access points can be viewed collectively as a single node. Assume traffic is in units of packets, and suppose each wireless node can transmit over at most one outgoing link per timeslot. The dashed links indicate time varying ON/OFF channels with i.i.d. ON probabilities of $1/2$ (so that a single packet can be transmitted when ON, and zero when OFF). Transmission costs over these links are 1 cent/packet, and




Fig. 1. A single commodity network where four different wireless sources desire to send data to any of the three wireline access points.

	SGP Profit	BB Profit	SGP \bar{U}	BB \bar{U}
User 1	0.36	1.36	—	—
User 2	1.13	1.70	—	—
User 3	0.72	1.08	—	—
User 4	1.31	1.05	—	—
Node 1	0.03	1.35	377	1533
Node 2	4.97	1.67	299	1211
Node 3	1.65	1.04	333	1574
Node 4	1.57	1.00	286	1438
Node 5	0.86	1.05	93	259
Node 6	0.02	0.59	132	782
Node 7	0.02	0.61	53	308
Total	12.63	12.51	1573	7104

Fig. 2. A table of simulation results for the single commodity network of Fig. 1, showing time average profit (in cents/slot) and queue backlog (in packets) for the SGP and Bang-Bang algorithms. $V = 50$ for SGP, $V = 500$ for BB. $R_{max} = 1$, $Q_{max} = g(1 + R_{max})$, $C_2 = C_5 = C_7 = 1$.

reception costs of all links are equal to 0.5 cent/packet. The solid links to the access points are always ON, and can transmit one packet per slot with transmission costs $C_2 = C_5 = C_7 = 1$. Suppose utility functions for each of the four users are given by $g(r) = 10 \log(1 + r)$. We simulate the SGP and Bang-Bang pricing algorithms over 10 million timeslots. We use $\Phi(x) = \Psi(x) = \log(1 + x)$ for the Bang-Bang implementation. Profit results are presented in Fig. 2. Note that nodes 5, 6, and 7 make positive profits by acting as pure relays. This shows that nodes can benefit from participating in the free market even if they do not desire their own personal communication.

Decreasing the V parameter in the Bang-Bang algorithm leads to less precisely balanced profits but also decreases congestion. Under SGP, nodes 6 and 7 support roughly 1/3 of the traffic from S4, but receive only 0.02 profit. In Fig. 3 we illustrate results for the case when C_2 and C_5 are increased to 3. In this case, it is apparent that the sum profit decreases, and the individual profits of nodes 2 and 5 significantly decrease under both SGP and BB, while the profit of nodes 6 and 7 either stay the same or increase.

B. An Ad-Hoc Mobile Network

Consider now an ad-hoc mobile network with 7 nodes that move about a network region that is partitioned into a 5×5 grid (see Fig. 4). Nodes 1 and 4 are stationary and stay in cell locations (4, 4) and (2, 2), respectively, for all time. Nodes 3, 5, 6, and 7 are fully mobile and take independent and uniform Markov random walks throughout the network. Specifically, every slot each fully mobile node independently decides with probability α to stay in its same cell, and else (with probability $1 - \alpha$) it decides to move to an adjacent cell either to the North, West, South, or East, equally likely in all 4 directions. If it decides to move in an infeasible direction, it stays in its same cell. Node 2 is *partially mobile*, in that it takes a Markov random walk that is restricted to only the two shaded cells in the bottom corner of Fig. 4. For simplicity, we assume that nodes can only communicate when they are in the same

	SGP Profit	BB Profit	SGP \bar{U}	BB \bar{U}
User 1	0.36	1.03	–	–
User 2	1.13	1.08	–	–
User 3	0.72	0.71	–	–
User 4	1.31	0.68	–	–
Node 1	0.03	1.01	377	1878
Node 2	2.97	1.05	299	1504
Node 3	1.04	0.66	333	1950
Node 4	0.96	0.63	286	1761
Node 5	0.07	0.64	154	924
Node 6	0.14	0.61	143	852
Node 7	0.03	0.62	53	309
Total	8.76	8.71	1646	9178

Fig. 3. A table of simulation results for the same single commodity network as Figs. 1 and 2, with the exception that C_2 and C_5 are increased to 3.

cell, that each cell can support at most one packet transmission per slot (consisting of 1.0 units of data), and that there is no inter-cell interference. The capacity of such a network when all nodes are fully mobile is derived in [33]. We assume a constant reception cost $C_r = 0.25$ and transmission cost of $C_t = 0.5$ units for each node. The α mobility probability is set to $\alpha = 1/2$.

We assume there are 3 sources, nodes 1, 2, and 3, and that each source node i desires to send a stream of data to a particular destination node j . The source-destination pairings are: $1 \rightarrow 4$, $2 \rightarrow 5$, and $3 \rightarrow 6$. Thus, this is a 3-commodity network, with commodities corresponding to the destination nodes 4, 5, and 6. Note that the stationary node 1 desires to send data to the other stationary node 4. Hence, the source-destination pair $1 \rightarrow 4$ requires mobile relays to support its communication. The other source-destination pairs can in principle exchange data without relays by waiting until they are in the same cell, but can considerably expand their throughput capabilities by using mobile relays. We assume that all nodes participate in the market and can act as relays of other data (including the stationary nodes 1 and 4). Note that node 7 is neither a source nor a destination, and hence it participates only to make a profit. For brevity, we simulate only the case when all nodes use the SGP algorithm.

It is important to note that this ad-hoc mobile network satisfies neither the i.i.d. channel assumption nor the orthogonal channel assumption. Indeed, the channel process between any node pair is effectively an ON/OFF process that depends on whether or not the nodes are in the same cell (which is not i.i.d. over slots due to the Markovian mobility). However, the SGP algorithm guarantees bounded queue backlog and non-negative profits for any stochastic channel process, including non-i.i.d. channels (recall Theorem 2 parts (a) and (b)). Further, while the maximum sum-profit result of Theorem 2 part (c) was proven for i.i.d. channels, the same result can be shown for channels modulated by any finite state ergodic Markov chain. The only difference is that the $O(1/V)$ term that bounds deviation from optimality is increased by a constant coefficient that depends on the mixing time of the Markov chain. This can be proven using a more detailed K -slot Lyapunov drift argument (see [30] for details on the technique), although the analysis is omitted for brevity.

Further, the requirement that each cell can support at most one transmission per slot violates the channel orthogonality assumption. This necessitates *transmission scheduling* in cases when there are multiple competing transmission possibilities in a given cell. However, it is simple to see that the proof of Theorem 2 holds exactly in this transmission scheduling scenario under the assumption that transmissions are chosen to maximize $\sum_n h_n(\tau)$ every slot, where $h_n(\tau)$ is defined in (23). Indeed, the orthogonality assumption was used only because it implies that maximizing this sum is equivalent to individually maximizing each $h_n(\tau)$ term, which can be done independently by each node n . However, in this cell-partitioned network scenario, maximizing the sum amounts to having each cell with two or more users select the transmitter-




Fig. 4. An ad-hoc mobile network under SGP.

	SGP Profit	SGP \bar{U}_i	SGP \bar{r}_i
User 1	0.030	—	0.105
User 2	0.038	—	0.112
User 3	0.045	—	0.138
Node 1	0.182	160.3	—
Node 2	0.269	178.0	—
Node 3	0.341	175.4	—
Node 4	0.026	70.5	—
Node 5	0.069	87.9	—
Node 6	0.058	91.0	—
Node 7	0.137	129.1	—
Total	1.196	892.2	0.354996

Fig. 5. A table of simulation results for SGP applied to the ad-hoc mobile network of Fig. 4, with $V = 20$. The average queue values \bar{U}_i correspond to the sum over all 3 queues within each node i .

receiver pair (n, b) with the largest differential price parameter $W_{nb}^*(\tau)$, and transmitting 1.0 units of the optimal commodity $c_{nb}^*(\tau)$ whenever $W_{nb}^*(\tau) - C_r - C_t > 0$. This transmission scheduling option has the following economic interpretation: Suppose that there is a “cell owner” that receives a portion of profits from each transaction in its region, and that arbitrates scheduling opportunities. The cell owner thus greedily selects the scheduling opportunity that maximizes instantaneous profit $W_{nb}^*(\tau) - C_r - C_t$ within the cell. This decision can alternatively be motivated with an auction strategy, under the assumption that each user bids for a transmission opportunity in proportion to its instantaneous profit $W_{nb}^*(\tau) - C_r - C_t$, and the highest bidder “wins.”

We simulate the network over 1 million timeslots, using utility functions $g(r) = 5 \log(1 + r)$ for each user 1, 2, 3, and $R_{max} = 1.0$. Note that this implies each queue has a worst case congestion of $5V + \delta_{max}$, where $\delta_{max} = 2.0$ in this scenario. Figs. 6 and 7 plot sum profit and average network congestion versus the V parameter, illustrating that sum profit converges to its maximum value as V is increased, with a corresponding linear tradeoff in average queue congestion. The “knee” of the curve appears roughly when $V = 20$, and hence we provide more detailed profit and queue backlog information for the $V = 20$ simulation in Fig. 5. Here we again make a distinction between the “user” at source node i and source node i itself (although the total profit of node i can also be considered as the sum of its user and node profits). Here we see that each of the three sources i achieves a throughput \bar{r}_i of at least 0.1 packet/slot, and all nodes and users make a non-negative profit.⁷

It is interesting to note that the destination nodes 4, 5, 6 make the least profit, as these nodes are mainly

⁷Note that the largest possible throughput between two mobile nodes would be only $1/25 = 0.04$ packets/slot if they did not use any relays (as this is the probability that both nodes are in the same cell). Hence, as all three sources achieve throughput larger than 0.1 packets/slot, we see that they are all actively utilizing relays.

interested in receiving their data and hence have the least opportunities for relaying data to others. Node 4 is stationary and hence makes the least profit of 4, 5, 6, as it is the least useful for relaying. Nodes 5 and 6 are both fully mobile, although node 5 makes more profit than 6 as the rate of data destined for 6 is larger than that of data destined for 5 (i.e., $r_3 > r_2$). The source nodes 1, 2, and 3 make the most profit, mostly coming from the users at these nodes. It is interesting to note that node 7 is neither a source nor a destination and hence is not interested in its own communication. However, it participates in the market by acting as a mobile relay, and makes a significant profit for these services. It makes a larger profit than the mobile nodes 5 and 6, as these destination nodes lose some relaying opportunities because they must receive their own data.

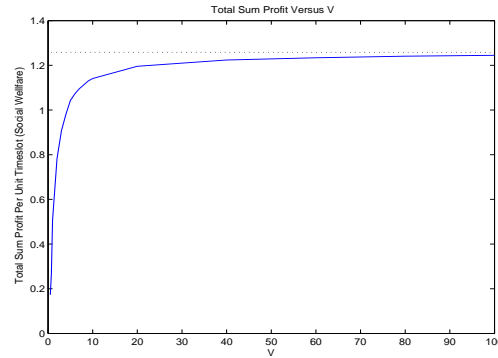


Fig. 6. Network profit versus V for the ad-hoc mobile network.

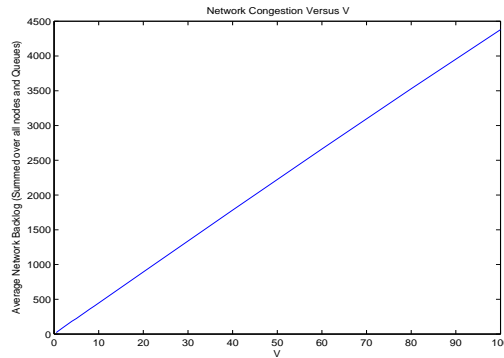


Fig. 7. Average network congestion (summed over all nodes and queues) versus V for the ad-hoc mobile network.

VI. CONCLUSION

This work addresses the problem where multiple self-interested wireless users must coordinate to support network communication. Two profit-aware algorithms were developed, one that maximizes sum profit while ensuring non-negative profits to all participants, and the other that achieves a more balanced profit sharing among users. The algorithms are based only on local message and data exchanges and do not require global knowledge of network topology or channel statistics. It was shown in both a static and ad-hoc mobile context that network participants can make a significant profit for their relaying services, even if they are not interested in their own personal communication. This work provides structured methods to coordinate dynamic markets, and can perhaps be used in other areas of network economics.

REFERENCES

- [1] M. J. Neely. Optimal pricing in a free market wireless network. *Proc. IEEE INFOCOM*, May 2007.
- [2] J. K. MacKie-Mason and H. R. Varian. Pricing congestible network resources. *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 7, September 1995.
- [3] F. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, vol. 8, pp. 33-37, 1997.
- [4] F.P. Kelly, A. Maulloo, and D. Tan. Rate control for communication networks: Shadow prices, proportional fairness, and stability. *Journ. of the Operational Res. Society*, 49, p.237-252, 1998.
- [5] S. H. Low and D. E. Lapsley. Optimization flow control, i: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, vol. 7(6): 861-75, Dec. 1999.
- [6] J. W. Lee, R. R. Mazumdar, and N. B. Shroff. Downlink power allocation for multi-class cdma wireless networks. *Proc. IEEE INFOCOM*, 2002.
- [7] M. J. Neely. *Dynamic Power Allocation and Routing for Satellite and Wireless Networks with Time Varying Channels*. PhD thesis, Massachusetts Institute of Technology, LIDS, 2003.
- [8] M. J. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. *Proc. IEEE INFOCOM*, March 2005.
- [9] L. Chen, S. H. Low, M. Chiang, and J. C. Doyle. Cross-layer congestion control, routing, and scheduling design in ad hoc wireless networks. *Proc. IEEE INFOCOM*, April 2006.
- [10] M. Andrews. Maximizing profit in overloaded networks. *Proc. IEEE INFOCOM*, March 2005.
- [11] D. Acemoglu, A. Ozdaglar, and R. Srikant. The marginal user principle for resource allocation in wireless networks. *Proc. of 43rd IEEE Conf. on Decision and Control*, Dec. 2004.
- [12] P. Marbach and R. Berry. Downlink resource allocation and pricing for wireless networks. *Proc. IEEE INFOCOM*, 2002.
- [13] T. Basar and R. Srikant. Revenue-maximizing pricing and capacity expansion in a many users regime. *Proc. IEEE INFOCOM*, June 2002.
- [14] L. He and J. Walrand. Pricing internet services with multiple providers. *Proc. of Allerton Conf. on Comm., Control, and Computing*, 2003.
- [15] L. He and J. Walrand. Pricing differentiated internet services. *Proc. IEEE INFOCOM*, March 2005.
- [16] H. Shen and T. Basar. Incentive-based pricing for network games with complete and incomplete information. *Proc. 11th International Symposium on Dynamic Games and Applications*, Dec. 2004.
- [17] I. Ch. Paschalidis and J. N. Tsitsiklis. Congestion-dependent pricing of network services. *IEEE/ACM Transactions on Networking*, vol. 8, no. 2, pp. 171-181, April 2000.
- [18] X. Lin and N. B. Shroff. Simplification of network dynamics in large systems. *IEEE/ACM Transactions on Networking*, vol. 13, no. 4, pp. 813-826, August 2005.
- [19] L. Buttyan and J.-P. Hubaux. Stimulating cooperation in self-organizing mobile ad hoc networks. *ACM/Kluwer Mobile Networks and Applications (MONET)*, vol. 8, no. 5, pp. 579-592, Oct. 2003.
- [20] J. Crowcroft, R. Gibbens, F. Kelly, and S. Ostring. Modeling incentives for collaboration in mobile ad-hoc networks. *presented at the 1st Int. Symp. Modeling and Optimization in Mobile, Ad-Hoc, and Wireless Networks (WiOpt '03)*, Sophia-Antipolis, France, March 2003.
- [21] L. Shang, R. P. Dick, and N. K. Jha. Desp: A distributed economics-based subcontracting protocol for computation distribution in power-aware mobile ad hoc networks. *IEEE Trans. on Mobile Computing*, vol. 3, no. 1, Jan.-March 2004.
- [22] P. Marbach and Y. Qui. Cooperation in wireless ad hoc networks: A market-based approach. *IEEE/ACM Transactions on Networking*, vol. 13, no. 6, Dec. 2005.
- [23] J. Huang, R. Berry, and M. Honig. Auction-based spectrum sharing. *ACM/Kluwer Journal of Mobile Networks and Applications (MONET) special issue on WiOpt'04*, to appear.
- [24] J. Sun, E. Modiano, and L. Zheng. Competitive fairness: Wireless channel allocation using an auction algorithm. *IEEE Journal on Selected Areas in Communications*, 2006.
- [25] Z. J. Ji, W. Yu, and K. J. R. Liu. An optimal dynamic pricing framework for autonomous mobile ad hoc networks. *Proc. IEEE INFOCOM*, April 2006.
- [26] Rahul Jain. *Efficient Market Mechanisms and Simulation-based Learning for Multi-Agent Systems*. PhD thesis, University of California, Berkeley, Fall 2004.
- [27] R. Maheswaran and T. Basar. *Auctions for divisible resources: Price functions, Nash equilibrium and decentralized update schemes*. Springer Verlag, pp. 87-102, 2002.
- [28] 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. 1936-1949, Dec. 1992.
- [29] M. J. Neely, E. Modiano, and C. E. Rohrs. Dynamic power allocation and routing for time varying wireless networks. *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 1, pp. 89-103, January 2005.
- [30] L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Foundations and Trends in Networking*, vol. 1, no. 1, pp. 1-149, 2006.
- [31] M. J. Neely. Energy optimal control for time varying wireless networks. *IEEE Transactions on Information Theory*, vol. 52, no. 7, July 2006.
- [32] D. P. Bertsekas and R. Gallager. *Data Networks*. New Jersey: Prentice-Hall, Inc., 1992.
- [33] M. J. Neely and E. Modiano. Capacity and delay tradeoffs for ad-hoc mobile networks. *IEEE Transactions on Information Theory*, vol. 51, no. 6, pp. 1917-1937, June 2005.