Optimizing Adaptive Modulation in Wireless Networks via Multi-Period Network Utility Maximization

Daniel O'Neill*, Ekine Akuiyibo*, Stephen Boyd* and Andrea J. Goldsmith*
*Department of Electrical Engineering,
Stanford University
Stanford CA 94305
Email: {dconeill, ekine, boyd, andrea}@stanford.edu

***Abstract*— We present a crosslayer technique to find and characterize optimal control policies for wireless networks operating at different time scales at the upper layer and physical layer. The technique can also be directly applied to networks carrying traffic with different time dependencies such as data or video. Our approach combines network utility maximization and adaptive modulation over an infinite discrete time horizon using a class of performance measures we call time smoothed utility functions. We describe the properties of optimal physical layer power and link rate policies and characterize optimal upper layer policies, which determine when packets should be injected into the network. We also characterize the behavior of optimal policies as different system parameters are used.**

I. Introduction

Adaptive modulation (AM) is a technique used to maximize spectral efficiency by adapting a link's transmitter power and rate under random channel conditions [1], [2], [3], [4], [5], [6]. We call this technique SE/AM for spectrally efficient adaptive modulation. SE/AM has proved useful in many applications, but it is an approach aimed at optimizing a physical layer metric only. It does not take into account the impact of upper layer protocols nor consider the characteristics of the traffic carried over the link.

To capture the effect of upper layer protocols, we combined network utility maximization (NUM) [7], [8], [9] and SE/AM to yield NUM/AM [10]. NUM/AM models the upper layer performance of data flows through a wireless network using utility functions. Different utility functions capture the behavior of different upper layer protocols (e.g. TCP.) Utility functions are functions of the rate at which data sources inject packets into the network. NUM/AM models multiple data flows traversing multiple links and yields optimal AM link power and link rate policies. These policies are optimal over the distribution of channel states. However these policies do not capture time dependent characteristics of the traffic carried by the network, and they implicitly assume that upper layer protocols respond on the same timescale as physical layer AM policies.

We extend NUM/AM to a broader class of utility functions that we call time smoothed utility functions. Time smoothed utility functions model the upper layer performance of data flows through a wireless network as functions of the time averaged rate at which packets are injected into the network. The time averaging serves two purposes. First, it captures differences in the characteristics of different data types, and second it reflects the observation that upper layer protocols often operate at longer time scales than those used by the physical layer. We call this extension multi-period NUM/AM (MPNUM/AM), since it averages over multiple time periods. MPNUM/AM is data source driven, with the data sources driving the demand for data throughput. This demand is supplied by the network at the expense of using transmitter power. MPNUM explicitly models the tension between the performance of upper layer network protocols and physical layer power costs. Time averaged flows are also considered in [11], [12] for buffered video and other transmissions.

Our results focus on describing and characterizing the optimal data source, link power and link rate policies for the wireless network. We formulate the problem for a general wireless network but, due to space limitations, consider only a single link with multiple data flows traversing it. Our analysis of optimal MPNUM/AM policies shows that these protocols are different than those of either SE/AM or NUM/AM. MPNUM/AM policies are functions of both the channel state and the average data source rate (network state.) Unlike SE/AM policies which have a fixed channel threshold below which no data will be transmitted, MPNUM policies will transmit even when channel conditions are poor if the average data source rate is too low. Unlike NUM/AM which always transmits, MPNUM/AM ceases to transmit if average source data rates are sufficiently high.

This paper is organized as follows: Section II describes the network model, and Section III presents the system performance metric. Section IV poses the problem as an infinite horizon average cost Markov decision process and describes a method of solution. Section V investigates a single link with multiple data flows. Section VI describes numerical examples, and Section VII summarizes our conclusions.

II. Model

We consider an infinite horizon discrete time model $t = 0, 1, \ldots$ of a multi-hop wireless network under flat fading. The network has $l = 1, \ldots, L$ links and $m = 1, \ldots, M$ flows. Multiple flows can traverse a given link, and a given flow can traverse multiple links from its source to its destination. Routing is described by matrix A, where $A_{lm} = 1$ if flow m traverses

[1]This work was supported in part by the DARPA ITMANET program under grant 1105741-1-TFIND and the AFOSR Complex Networks Grant FA9550-08-1-0480

link l and $A_{lm} = 0$ otherwise. The channel gain matrix G^t describes the matrix of channel gains at time t, where G_{ij}^t is the gain from the transmitter on link j to the receiver on link i. The gain matrix is random and models flat fading. We assume the underlying process is Markov. The network is assumed to have knowledge of current channel gains.

The network can adapt its link transmitter powers, link rates, and the instantaneous rate at which packets associated with a given flow enter the network. Link l transmits at power ϕ_l^t at time t. The vector of transmitter powers ϕ^t is determined by an AM power policy which is assumed to be a function of the channel gain matrix G^t and the vector of average flow rates r^t, which is described below. For compactness of notation, the value of the power policy is written as $\phi^t = \phi(G^t, r^t)$. Link l transmits at rate R_l^t at time t. The rate policy vector is a function of the channel gain matrix and vector of link powers and is written as $R^t = R(G^t, \phi^t)$.

Each source injects packets into the network at instantaneous flow rate u_m^t. The vector of instantaneous source flow rates u^t is a function of G^t and the time averaged flow rates r^t. The instantaneous source flow rate policy vector is written as $u^t = u(r^t, G^t)$.

At any time t the total of all flow rates traversing a link must be less than the transmission rate of that link. This can be written as

$$Au^t \leq R^t. \tag{1}$$

Upper layer protocols are modeled as utility functions [13], [14]. TCP in particular [15] has been modeled in this way. Associated with each flow is a utility function $\mathbf{U}_m(r_m^t)$ which measures the upper layer performance of averaged flow m at time t. Different flows may have different utility functions, reflecting the use of different protocols. Utility functions are increasing strictly concave functions. Strictly concave utility functions exhibit diminishing returns with rate, that is, as rate increases the incremental utility grows by smaller amounts. Video protocols often exhibit this property. We consider utility functions of the form

$$\mathbf{U}(r) = \begin{cases} \frac{r^{1-\alpha}}{1-\alpha} & 0 < \alpha < 1 \\ \log r & \alpha = 1 \end{cases} \tag{2}$$

The parameter α corresponds to different properties of the utility function. For $\alpha = 1$ the utility function has the property of proportional fairness.

The variable r exponentially averages the instantaneous flow rates

$$r^{t+1} = \theta r^t + (I - \theta)u^t, \tag{3}$$

where I is the identity matrix and θ is a diagonal matrix with $0 \leq \theta < I$. We call r^t the state of the system. Averaging the flow rates reflects the demands of different types of traffic. When $\theta = 0$ each period is evaluated independently. This models traffic that is delay sensitive or where packets can't be shifted between time periods. Voice traffic, with the appropriate utility function, can be modeled in this manner. For file transfer, packets can be shifted between periods, with the average rate r a more important metric than the instantaneous rate. In this case $\theta \approx 1$ may be appropriate. For buffered video traffic, short term averages may be most appropriate and an intermediate value of θ can be used. More complex averaging schemes can be used, but for clarity we use this simple method of averaging.

III. System Performance

Our network performance metric is the average net benefit associated with a set of policies. Net benefit is the difference between the upper layer performance and physical layer power cost of the system. In a single time period this is expressed as

$$E\left[\sum_m \mathbf{U}_m(r_m^t) - \nu^T \phi(G^t, r^t)\right], \tag{4}$$

where E is the expectation operator over channel states. The parameter $\nu \geq 0$ determines the tradeoff between average power and average utility and can be thought of as the power cost of performance in units of performance per Watt.

Network performance is defined as the time average of the per-period performance over an infinite time horizon. This is written as

$$J_{u,\phi}(r_0) = \lim_{N\to\infty} \frac{1}{N} E\left[\sum_{t=0}^{N-1} (\sum_m \mathbf{U}_m(r_m^t) - \nu^T \phi(G^t, r^t))\right], \tag{5}$$

where $r_0 > 0$ is the vector of initial flow rates at each source and the expectation is over $\{G^t\}_{t=0}^{\infty}$. The variable $J_{u,\phi}(r_0)$ is the average net benefit of using policies u and ϕ. Different instantaneous source flow policies u or AM power policies ϕ can result in different average benefit.

The time averaged cost function emphasizes the long term behavior of the system and does not discount its future behavior. We use this formulation as opposed to a discounted performance metric on pragmatic grounds; without prior knowledge, each byte of a file or voice message is of equal importance and there is little justification in discounting later packets as inherently less important. As such (5) seeks to capture the long term cost/benefit of using an instantaneous flow function u. Transient network behaviors are "averaged out" in (5).

The objective is to describe the instantaneous source flow rate policy u and AM power policy ϕ that maximize (5). That is, we seek stationary functions u^* and ϕ^* such that

$$J_{u^*,\phi^*}(r_0) = \sup_{\pi\in\Pi} J_\pi(r_0), \tag{6}$$

where π is any time sequence of stationary functions contained in a defined "admissible" class Π. In this paper Π is the class of functions of the form $\gamma(r, G)$ and $\beta(G, r)$.

Equations (5) and (6) describe an average cost Markov Decision Process, MDP, [16], [17]. The system state variable is r. The instantaneous source flow rate policy u and AM power policy ϕ are stationary functions of the channel state and system state. Note that we are using the term policy to refer to both the function and (more properly) to the sequence of using the function repeatedly over the infinite horizon. As intuition would suggest, the average cost associated with the optimal policies is independent of the initial state r_0 and can be written as the constant J_{u^*,ϕ^*}.

IV. Method of Solution

Bellman's equation describes the optimality condition for an average cost MDP. The wireless network uses knowledge of its state r and channel conditions G to adapt to changing conditions. This assumption allows us to use the post-decision form of the Bellman equation [18]. Policies u^* and ϕ^* are optimal if one can find constant J_{u^*,ϕ^*} and function $V(r)$ such that

$$J_{u^*,\phi^*} + V(r) \quad = \quad E\{\max_{u,\phi\geq 0}(\sum_m \mathbf{U}_m(r_m) - \nu^T\phi + V(\theta r + (I-\theta)u)\}. \tag{7}$$

The function $V(r)$ is called the relative value function and captures the deviation from average system performance when the system is in state r. Unfortunately equation (7) is very difficult to solve analytically [17] and numerical methods must be used. Consequently, we characterize the properties of optimal policies, describing how policies change as system parameters are modified and then numerically solve for optimal policies.

Our approach focuses on a convex optimization problem associated with (7). Formally, the maximization in (7) is over positive functions u and ϕ. However in post-decision form, the maximization is inside the expectation operator, allowing values of the optimal policies to be calculated as variables from an associated real valued optimization problem. In particular, for given values of r and G the optimal values of the policies $u(r,G)$ and $\phi(G,r)$ can be expressed as

$$[u(r,G),\phi(G,r)] \quad = \quad \operatorname*{argmax}_{\hat{u},\hat{\phi}\geq 0}\{\sum_m \mathbf{U}_m(r_m) - \nu^T\hat{\phi} + V(\theta r + (I-\theta)\hat{u})\} \tag{8}$$

where the optimization is over the positive real variables $\hat{u}$ and $\hat{\phi}$.

V. A Wireless Link

In this section we characterize the optimal instantaneous source flow rate u and AM link power policies ϕ for a single link. We first consider a single flow to illustrate the approach and then analyze a single link with multiple flows. In both cases we assume MQAM modulation and that the link rate function is

$$R^t = \log\left(1 + K\frac{G^t\phi^t}{N}\right), \tag{9}$$

where $K = \frac{-1}{\log(BER)}$, and BER is the target bit error rate [19]. For simplicity we normalize $N = 1$.

In the single flow case, (1) simplifies to $u^t = R^t$. Thus the link power policy ϕ is determined by u and can be expressed as

$$\phi^t = \frac{\exp^{u^t} - 1}{KG^t}, \tag{10}$$

and only a single policy needs to be found. In this form the optimal instantaneous source flow policy drives the optimal link power policy and link rate. That is, u represents the demand for throughput, and the optimal policies supply this demand. The associated optimality equation (8) becomes

$$u(r,G) \quad = \quad \operatorname*{argmax}_{\hat{u}\geq 0}\{\mathbf{U}(r) - \nu\frac{\exp^{\hat{u}}-1}{KG} + V(\theta r + (1-\theta)\hat{u})\}. \tag{11}$$

The per-period network performance is strictly concave and increasing in r and consequently it can be shown through induction that the relative value function $V(r)$ is also strictly concave and increasing. Thus, for each G and r, (11) can be viewed as a constrained concave optimization problem with a global maximum.

Letting $y = \theta r + (1-\theta)u$, (11) becomes

$$y^*(r,G) = \operatorname*{argmax}_{\hat{y}\geq\theta r}\left\{\mathbf{U}(r) - \nu\frac{\exp(\frac{\hat{y}-\theta r}{(1-\theta)}) - 1}{KG} + V(\hat{y})\right\} \tag{12}$$

where y^* is the optimal value at channel state G and system state r. The optimal instantaneous source rate policy then has the form

$$u^*(r,G) = \begin{cases} [y^* - \theta r][1-\theta]^{-1} & y^* \geq \theta r \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

At a given system state r, the system will try to maintain a target average flow rate y^* for each channel condition G. If the average flow rate is below y^* the link will transmit at a rate sufficient to return the average flow rate to y^*. If the average flow rate exceeds y^*, then nothing will be transmitted that period and in subsequent periods until the average rate declines below this threshold. The rate of decline is determined by θ, with $k = \frac{\log(y^*/r)}{\log(\theta)}$ time steps required to the next transmission for a constant G.

Properties of $y^*(r,G)$ can be deduced from (12) using Topkis' Monotonicity Theorem [20]. Topkis' Theorem relates the properties of supermodularity and monotonicity. A function $f(x,\beta)$ is supermodular in (x,β) if for any $x' > x$ $f(x',\beta) - f(x,\beta)$ is nondecreasing in β. If we think of x as the argument and β as a parameter, a supermodular function is nondecreasing for an increasing parameter β. Topkis' Theorem states that if f is supermodular in (x,β), then

$$x^*(\beta) = \operatorname*{argmax}_{x\in D} f(x,\beta) \tag{14}$$

is nondecreasing in β, where D is the set of feasible x values.

Equation (11) is supermodular in G when the other parameters are held constant. Consequently y^* is nondecreasing in G, matching the intuition that the optimal policy seeks to boost the target flow rates associated with better channel conditions. Equation (11) is submodular in ν. Thus, the average power is nonincreasing as ν is increased. This matches the intuition that the optimal policy will decrease the average power used as it becomes more costly.

The averaging factor θ controls the coupling between discrete time periods. For a fixed channel state, the source will inject packets sufficient to return the system to the target flow rate $y^*(r,G)$. Equation (12) is supermodular in θ. Thus, y^* is nondecreasing in θ. A special case is $\theta = 0$, since each period is independent of every other period. In this case the optimal policy is described in [10] and the network will transmit in every time period. As θ gets larger the impact of the link transmitting is felt to a greater degree in future periods, thus increasing average utility at no incremental power.

The optimal MPNUM/AM policies are different than those of either SE/AM or NUM/AM. The MPNUM/AM power policy is a function of the channel state G and system state r. Unlike SE/AM policies which have a fixed channel threshold below which no data will be transmitted, MPNUM policies will transmit even when channel conditions are poor if the average data source rate is too low. Unlike NUM/AM which always transmits, MPNUM/AM ceases to transmit if average source data rates are sufficiently high. In general MPNUM/AM will transmit whenever $y^*(r,G) \geq \theta r$ and like SE/AM and NUM/AM will increase transmitter power with improving channel conditions.

A. Multiple Flows

Multiple data flows, $M \geq 2$, across a single link are modeled in a similar manner. The link rate bounds the sum of the instantaneous flows and (1) becomes $\mathbf{1}^T u^t = R^t$. As in the single flow case, the instantaneous source rate policies determine the AM power policy yielding

$$\phi^t = \frac{\exp(\mathbf{1}^T u^t) - 1}{KG^t}. \tag{15}$$

The associated optimality equation is

$$\begin{aligned} u(G,r) \quad = \quad & \underset{\hat{u}\geq 0}{\operatorname{argmax}}\{\mathbf{U}(r) - v\tfrac{\exp(\mathbf{1}^T\hat{u})-1}{KG} + \\ & \qquad V(\theta r + (I-\theta)\hat{u})\}. \end{aligned} \tag{16}$$

Analysis is similar to the single flow case and yields optimal policies with more complex instantaneous flow rate thresholds. Writing $y(r,G) = \theta r + (I-\theta)u$ and defining y^* as the optimal value of the related unconstrained problem yields the optimal policy

$$u^*(r,G) = \begin{cases} [y^* - \theta r][I-\theta]^{-1} & y^* \geq \theta r \\ \hat{u}^* & \text{otherwise} \end{cases} \tag{17}$$

where the inequalities are element wise. Each instantaneous flow is associated with a threshold $y^*_m(r,G)$. As in the single flow case, when $y^* \geq \theta r$ the demands of the instantaneous source vector determine the throughput supplied by the network. When $y^* \leq \theta r$ for one or more flows, the optimal instantaneous flow rates $\hat{u}^*$ for those flows may still be nonzero, depending on the particular set of utility functions selected. The quantity $\hat{u}^*$ is computed from (16).

VI. Numerical Results

In this section we present results for a discrete approximation for the case of a single flow over a single link. The state space and policy space are partitioned and relative value function for the discrete problem computed iteratively from the Bellman equation (7) using relative value iteration. Although this approach converges for the problems shown here, convergence is not, in general, guaranteed [17]. The

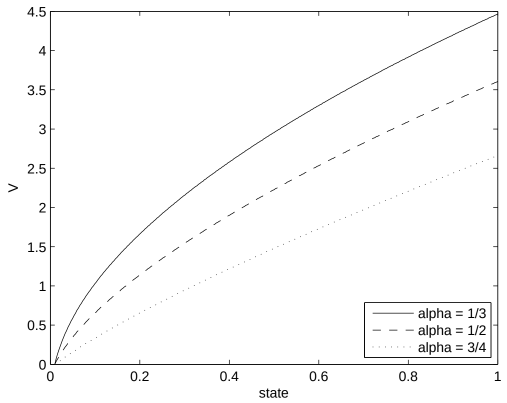

Fig. 1. Relative Value Function

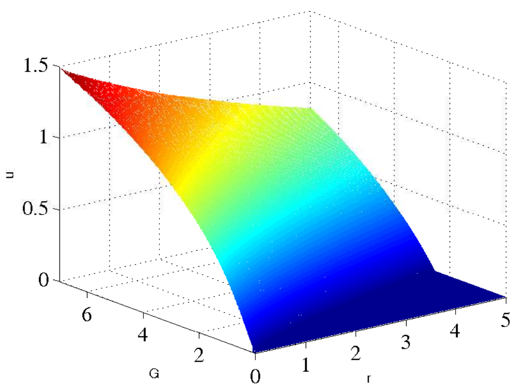

Fig. 2. Flow Policy

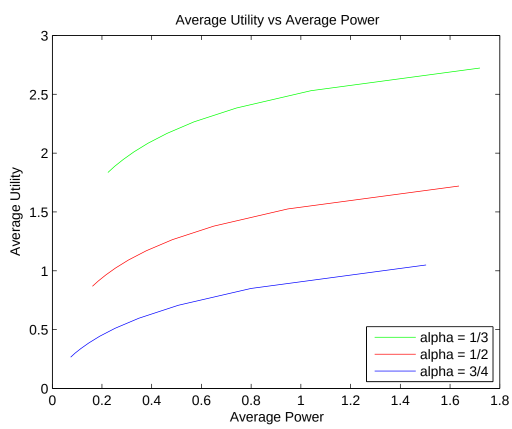

Fig. 3. Ave Utility vs. Average Power

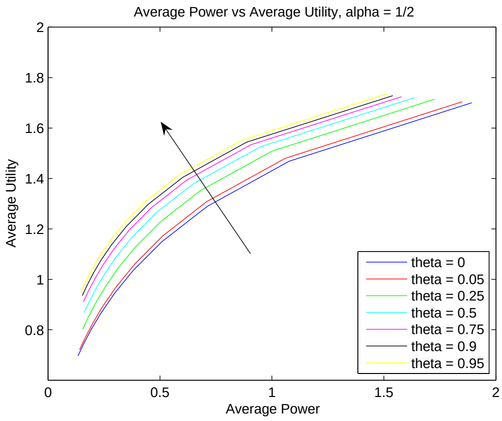

Fig. 4. Ave Utility vs Ave Power

transition probabilities used are based on IID Rayleigh fading of the elements of G.

The relative value iteration algorithm recursively computes the optimal values of the average cost J_{u^*} and relative value function $V(r)$ for the discretized problem. Let $V_i^{(k)}$ be the kth iterate of V_i, $\hat{V}_i^{(k)}$ is the kth iterate of $\mathscr{T}V_i$, and $J^{(k)}$ is the kth iterate of J. The value iteration algorithm is as follows:

for $i = 1, \ldots, N$
initialize $V_i^{(0)} = 0$.
repeat

1. $\hat{V}_i^{(k)} := 0$.
2. $\hat{V}_i^{(k)} := \mathscr{T}V_i^{(k)}$.
3. *Update.* $J^{(k)} := \hat{V}_1^{(k)}$.
4. *Stopping criterion.* **quit** when $J^{(k)}$ converges.
5. *Update.* $V_i^{(k+1)} := \hat{V}_i^{(k)} - \hat{V}_1^{(k)}$.

where $\mathscr{T}$ is the Bellman operator [17].

Figure 1 shows that the value function is strictly concave and increasing for three different values of α. Figure 2 shows the optimal policy for the case of a utility function with $\alpha = \frac{1}{2}$. As can be seen, the policy is a function of both r and G. Transmitter power cut-off occurs for larger values of G as the system state r increases as expected. Figure 3 plots average power versus average utility for different values of ν. The three curves correspond to three different values of α. Figure 4 shows average power versus average utility as ν is varied for seven different values of θ, the averaging parameter. The curves are produced by fixing θ and calculating average power and average utility for a range of values for ν.

VII. Conclusion

We extend NUM/AM to a broader class of utility functions that we call time smoothed utility functions. Time smoothed utility functions reflect differences in the time dependent characteristics of different types of traffic and also the different time scales used by different layers of the the protocol stack. We term this extension as MPNUM/AM since it averages over multiple periods and is modeled as an infinite horizon average cost Markov decision process. MPNUM/AM captures the tension between the demand for upper layer performance and the power cost of supplying it.

The MPNUM/AM power and link rate policies are functions of both channel conditions and network state and are different than either SE/AM policies or NUM/AM policies. MPNUM policies have no fixed transmission threshold, but will transmit even when channel conditions are poor if the average flow rate is too low. Unlike NUM/AM which always transmits, MPNUM/AM ceases to transmit if average source data rates are sufficiently high. We characterize the behavior of the optimal MPNUM policies as different system parameters are used.

References

[1] R. Hoefel and C. de Almeida, "Performance of ieee 802.11-based networks with link level adaptive techniques," *Vehicular Technology Conference, VTC2004-Fall. 2004 IEEE 60th*, vol. 2, 2004.

[2] F. Peng, J. Zhang, and W. E. Ryan, "Adaptive modulation and coding for ieee 802.11n," *IEEE Wireless Communications and Networking Conference, 2007.WCNC*, 2007.

[3] K.-B. Song, A. S. T. C. Ekbal, and J. Cioffi, "Adaptive modulation and coding (amc) for bit-interleaved coded ofdm (bic-ofdm)," *IEEE Communications, 2004 IEEE International Conference on*, vol. 6, no. 11, 2004.

[4] X. Qiu and K. Chawla, "On the performance of adaptive modulation in cellular systems," *IEEE Trans. Comm.*, pp. 884–895, June 1999.

[5] A. Goldsmith, "The capacity of downlink fading channels with variable rate and power," *Vehicular Technology, IEEE Transactions on*, vol. 46, no. 3, 1997.

[6] S. Chung and A. J. Goldsmith, "Degrees of freedom in adaptive modulation: A unified approach," *IEEE Trans. Commun.*, vol. 49, 2001.

[7] J.-W. Lee, M. Chiang, and A. R. Calderbank, "Price-based distributed algorithms for rate-reliability tradeoff in network utility maximization," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 5, 2006.

[8] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "'layering as optimization decomposition: A mathematical theory of network architectures'," *Proceedings of the IEEE*, vol. 95, pp. 255–312, January 2007.

[9] L. Georgiadis, M. Neely, and L. Tassiulas, *Resource Allocation and Cross-Layer Control in Wireless Networks*. Hanover, MA: NOW Publishers, 2006.

[10] D. O. ONeill, A. J. Goldsmith, and S. Boyd, "Optimizing adaptive modulation in wireless networks via utility maximization," *ICC 2008*, 2008.

[11] F. Fu and M. van der Schaar, "A new systematic framework for autonomous cross-layer optimization," *IEEE Trans. Veh. Tech.*, pp. 1887–1903, May 2009.

[12] ——, "Decomposition principles and online learning in cross-layer optimization for delay-sensitive applications," *UCLA Technical Report*, 2009.

[13] R. Srikant, Ed., *The Mathematics of Internet Congestion Control*. Boston: Birkhauser, 2003.

[14] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communications networks: Shadow prices, proportional fairness and stability," *Journal of Operations Research Society*, vol. 49, pp. 237–252, Nov. 1998.

[15] S. Low, L. Peterson, and L. Wang, "Understanding vegas: A duality model," *Journal of ACM*, vol. 49(2), pp. 207–235, Mar. 2002.

[16] S. Meyn, *Control Techniques for Complex Networks*. New York, NY: Cambridge University Press, 2008.

[17] D. Bertsekas, *Dynamic Programming and Optimal Control*. Massachusetts: Athena Scientific, 2005.

[18] W. Powell, *Approximate Dynamic Programming*. Hoboken, NJ: John Wiley and Sons, 2007.

[19] G. J. Foschini and J. Salz, "Digital communications over fading radio channels," *Bell Syst. Tech. J.*, pp. 429–456, Feb. 1983.

[20] D. Topkis, *Supermodularity and Complemantarity*. Princeton: Princeton University Press, 1998.