Threshold Structure of Channel Aware Distributed Scheduling in Ad-Hoc Networks: An Optimal Stopping View

Junshan Zhang
Dept. of EE, Arizona State University, Tempe, AZ, 85287

Abstract— **As evidenced by measurement data, channel fading and co-channel interference occur on the same time scales, and it is therefore difficult to determine if packet losses are due to interference change or channel variation. The coupling between the timescales of fading and interference at the MAC layer calls for a unified PHY/MAC design. Using optimal stopping theory, we first devise channel aware distributed scheduling to exploit rich PHY/MAC diversities in single-hop ad-hoc networks, for a variety of PHY-layer models. We show that the optimal channel aware distributed scheduling algorithms have threshold structures, and hence are amenable to easy implementation. We then generalize the study to multi-hop wireless networks, and discuss further open issues such as the delay performance of distributed scheduling.**

Index Terms— Threshold Policy, Distributed Scheduling, Ad-Hoc Networks, Optimal Stopping.

I. Introduction

Wireless network design faces two unique challenges in wireless communications, namely co-channel interference and time varying channel fading, and the combination of the two may result in a higher order of packet losses in wireless networks. Co-channel interference is due to other concurrent transmissions in the neighborhood. Channel fading is the time variation of the wireless channel, consisting of two effects: large-scale path loss and shadowing effects that cause the signal to attenuate with distance; and multipath scattering effects that result in delayed copies of the signal adding up constructively or destructively at the receiver.

The traditional approach for wireless network design intends to separate link losses caused by fading from those by interference. That is, the PHY layer addresses the issues of fading, while the MAC layer addresses that of contention. This hope for separation of point-to-point link reliability and multiple access functionality between the PHY and MAC layers relies on the implicit assumption that the PHY layer works perfectly and hides fading from MAC. However, as shown in the measurement data [1] [2], channel fading and co-channel interference occur on the same time scales, and it is therefore difficult to determine if packet losses are due to MAC-layer interference variation or channel variation. The coupling between the timescales of fading and MAC calls for a unified PHY/MAC design, and there is an urgent need to develop a rigorous foundation of PHY-aware scheduling.

Distributed scheduling has received much attention over the past decade. In particular, since the seminal work [3] on throughput maximization for constrained queueing systems, there has recently been a surge of interest in devising distributed scheduling for multi-hop wireless networks. The complexity levels of distributed wireless scheduling have a wide spectrum, ranging from constant time complexity (e.g., decentralized pick-and-compare algorithms) to exponential complexity (e.g., max-weight scheduling) (see [4] and the references therein). Clearly, the overhead due to demanding message passing, if not addressed carefully, can drive the effective throughput to zero. Moreover, most existing distributed scheduling algorithms are not *channel-aware* in the sense that they do not exploit channel variations. Furthermore, from a practical perspective, it is preferable to develop simple-to-implement (particularly threshold-based) scheduling.

Channel aware opportunistic scheduling was first developed for downlink transmissions in cellular wireless networks (see, e.g., [5], [6], [7], [8], [9], [10], [11]), assuming that the centralized scheduler has knowledge of the instantaneous channel conditions for all links. More recently, channel aware random access has been investigated for the uplink transmissions in a many-to-one network [12], [13], where the channel probing can be realized by broadcasting pilot signals from the base station.

In this paper, we explore the threshold structure of channel aware distributed scheduling in wireless ad-hoc networks. Most existing studies on channel-aware scheduling require centralized scheduling, and little work has been done on developing distributed scheduling to exploit rich PHY/MAC diversities for ad

This research is supported in part by Office of Naval Research through the grant N00014-05-1-0636 and National Science Foundation through the grants ANI-0238550 and CNS-0721820.

hoc networks. This is perhaps due to the challenge that the distributed nature of ad hoc communications implies that each link has little knowledge of other links' channel conditions.

II. Channel-Aware Distributed Scheduling for Exploiting PHY/MAC Diversities

Consider a single-hop ad-hoc network, where links contend for the same channel using random access. We assume that after a successful probing, the channel condition of the successful link is measured. Due to channel fading, the link condition corresponding to this successful channel probing can be either good or poor. In the latter case, data packets have to be transmitted at low rates. A plausible alternative is to let this link give up this opportunity, and allow all the links re-contend for the channel, in the hope that some link with a better channel condition can transmit after the re-contention. Intuitively speaking, it is likely that after further probing, the channel can be taken by a link with a better channel condition, resulting in possible higher throughput. In this way, the multiuser diversity across links and the time diversity across slots can be exploited in a joint opportunistic manner. Fig. 1 depicts a sample realization with N rounds of channel probing, followed by one data transmission.

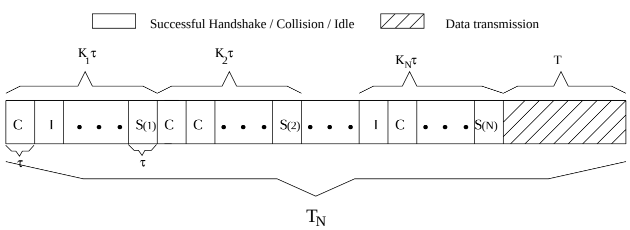

Fig. 1. A sample realization of channel probing and data transmission

Specifically, suppose there are M transmitter nodes, each with possibly multiple intended receivers. We assume that each transmitter node m contends for the channel using random access with probability p_m, $m = 1, \ldots, M$. In this section, we assume a collision model for channel contention, where a channel contention of a node is said to be successful if no other nodes in the same neighborhood transmit at the same time. Accordingly, the overall successful contention probability p_s is given by $\sum_{m=1}^{M} p_{s,m}$ with $p_{s,m} = p_m \prod_{i \neq m}(1-p_i)$; and the number of slots (denoted as K) needed to accomplish a successful channel contention is a Geometric random variable, i.e., $K \sim Geometric(p_s)$. Let $s(n)$ denote the successful link in the n-th round of channel probing, and $R_{n,s(n)}$ denote the corresponding transmission rate. We assume that $R_{n,s(n)}$ remains constant for a duration of T, where T is the data transmission duration and is no greater than the channel coherence time. Let $\tau < T$ denote the duration of mini-slot for channel contention.

It is clear that there is a fundamental *tradeoff* between the throughput gain from better channel conditions and the cost for further channel probing. The desired tradeoff boils down to judiciously choosing the optimal stopping rule for channel probing, in order to maximize the throughput. In what follows, we obtain a systematic characterization of this tradeoff by appealing to optimal stopping theory [14], [15].

A. Threshold Structure for The Single-Receiver Model

First consider the model where each transmitter is associated with one receiver only. We treat channel-aware distributed scheduling, as a team game in which all links collaborate to maximize the overall network throughput. For convenience, let $R_{(n)}$ denote the rate corresponding to the n-th round successful channel probing, i.e., $R_{(n)} = R_{n,s(n)}$. Without loss of generality, we assume that the second moment of $R_{(n)}$ exists.

Building on optimal stopping theory, we cast the problem as *maximizing the rate of return*, where the rate of return is the expected throughput [15], and a key step here is to characterize the optimal stopping rule N^* and the optimal throughput x^*, i.e.,

$$N^* \triangleq \arg\max_{N \in Q} \frac{E[R_{(N)}T]}{E[T_N]}, \quad x^* \triangleq \sup_{N \in Q} \frac{E[R_{(N)}T]}{E[T_N]}. \quad (1)$$

where T_N is the total duration for probing and data transmission, and

$$Q \triangleq \{N : N \geq 1, E[T_N] < \infty\}. \quad (2)$$

For the network model with homogenous links, i.e., all links have the same channel statistics with the same distribution $F_m(r)$. We have the following proposition [16].

Proposition 2.1: a) The optimal stopping rule N^* for channel-aware distributed scheduling exists, and is given by

$$N^* = \min\{n \geq 1 : R_{(n)} \geq x^*\}. \quad (3)$$

b) The maximum throughput x^* is an optimal threshold, and is the unique solution to

$$E(R_{(n)} - x)^+ = \frac{x\tau}{p_s T}. \quad (4)$$

In practical scenarios, it is likely that different links may have different channel statistics. As a result, if $s(n+$

$1) \neq s(n)$, $R_{n,s(n)}$ and $R_{n+1,s(n+1)}$ may follow different distributions. The following proposition summarizes the optimal threshold policy for this case.

Proposition 2.2: The optimal stopping rule N^* for channel-aware distributed scheduling exists, and the maximum throughput x^* in the heterogeneous case is an optimal threshold, and is the unique solution to the following equation:

$$x = \frac{\sum_{m=1}^{M} p_{s,m} \int_x^\infty r dF_m(r)}{\delta + \sum_{m=1}^{M} p_{s,m} \left(1 - F_m(x)\right)}. \tag{5}$$

Remarks: 1) Proposition 2.1 reveals that the optimal stopping rule N^* for channel-aware distributed scheduling is a pure threshold policy, and the stopping decision can be made based on the current rate only. (For the discrete rate case, we treat the thresholds in between two adjacent quantization levels "effectively" the same). Accordingly, the optimal channel probing and scheduling strategy takes the following simple form: If the successful link discovers that the current rate $R_{(n)}$ is higher than the threshold x^*, it transmits the data with rate $R_{(n)}$; otherwise, it skips this transmission opportunity, and then the links re-contend.

2) For the heterogeneous case, *a priori*, it is not clear that different links would have different thresholds or not since their channel statistics are different. However, Proposition 2.2 indicates that in the optimal strategy the threshold is the same for all the links . Our intuition is as follows: When all the links have the same threshold, links with better channel conditions would have a higher likelihood to transmit accordingly.

B. Threshold Structure for The Multi-Receiver Model

Consider now each transmitter is associated with multiple intended receivers. Different from the model where each transmitter is associated with one single receiver only, the probing in the multi-receivers case takes place in two phases: 1) In phase I, all transmitters contend for the channel using random access to reserve the channel, and the probing in this phase to accomplish a successful channel contention takes a *random duration of* $K\tau$; and 2) In phase II, subsequent probings are carried out to estimate the channel conditions from the successful transmitter in phase I to its intended receivers, according to specific probing mechanisms, and for each receiver the probing for channel condition incurs a *constant duration of* τ.

There are many possible probing mechanisms, and here we consider the sequential probing without recall (SPWOR) mechanism only. In SPWOR, after a successful contention, the transmitter probes its receivers sequentially, and stops the probing process once it probes a "good" channel, followed by data transmission to the most recently probed receiver. (See [17] for a detailed study of the random selection (RS) mechanism, the exhaustive sequential probing with recall (ESPWR) mechanism, and the sequential probing with recall (SPWR) mechanism).

Let $R_{n,t(n),j}$ denote the rate of receiver j in the nth round of channel probing. We have the following result [17].

Proposition 2.3: (**Multi-Receiver model**) a) Suppose that sequential probing without recall (SPWOR) is used for channel probing. Then the optimal stopping rule for distributed scheduling is given as follows:

$$N^*_{SPWOR} = \min\{\kappa \geq 1 : R_{n,t(n),j} \geq \theta^*_j, \\ \text{where } n = \lceil \frac{\kappa}{L} \rceil, j = \mod(\kappa - 1, L)\}, \tag{6}$$

and the thresholds $\{\theta^*_j\}$ are determined by

$$\theta^*_j = x^*_{SPWOR} + v^*_{j+1}, \forall\ j = 0, 1, \ldots, L-1, \tag{7}$$

b) The maximum network throughput x^*_{SPWOR} is the unique solution to the following fixed point equation:

$$E[\max(R - x, v^*_1(x))] - \frac{x\delta}{p_s} = 0, \tag{8}$$

where R is a random variable with distribution $F_R(r)$, and $\{v^*_j(x)\}$ are defined (in a backward order) as follows:

$$v^*_L(x) \triangleq 0, \tag{9}$$

$$v^*_j(x) \triangleq E[\max(R - x, v^*_{j+1}(x))] - x\delta, \\ \forall\ j = L-1, L-2, \ldots, 1. \tag{10}$$

c) $v^*_j \triangleq v^*_j(x^*_{SPWOR}), \forall\ j = 1, 2, \ldots, L$.

Remarks: Proposition 2.3 reveals that the optimal scheduling policy corresponding to SPWOR probing exhibits a multi-stage threshold structure. Furthermore, observe that the optimal thresholds given by (7) only depends on the number of receivers that the transmitter has probed, indicating that the optimal stopping rule in (6) is amenable to easy distributed implementation.

As expected, the optimal thresholds at earlier-probed receivers are larger than that at later-probed receivers, i.e., $\theta^*_i \geq \theta^*_j, \forall\ i \leq j$. Intuitively speaking, at receiver i, more receivers (i.e., $L - i - 1$ remaining receivers) are available for further probing (and can be possibly utilized), compared to at receiver j. Indeed, it can be shown that the thresholds $\{\theta^*_j, \forall\ j = 0, 1, \ldots, L-1\}$ in (7) monotonically decrease, i.e., $\theta^*_0 \geq \theta^*_1 \geq \cdots \geq \theta^*_{L-1}$.

III. Channel-Aware Distributed Scheduling With Physical Interference Model

Observe that a key assumption in the above studies is the collision model, which assumes that a channel contention of a link is successful if no other links transmit at the same time. However, with the advent of new signal processing techniques, such as multiuser detection, spread spectrum, and space-time processing, it is possible to simultaneously decode multiple packets even when "collision" happens. These new techniques at the PHY layer offer the capability of *Multipacket reception* (MPR), and we call the scheduling in this context "PHY-aware distributed scheduling".

Devising PHY-aware distributed scheduling with MPR is in general difficult. One unique challenge in exploiting MPR is that at each transmission, multiple links can transmit successfully through one common channel, and furthermore, each link has to make the decision to transmit or not based on local information only, because links involved in the transmission have no knowledge of the instantaneous information of other links. Furthermore, the number of transmitting links is random, and heavily depends on the contention probability of each link. Roughly speaking, a large contention probability would increase the number of probing links and thus incur strong cochannel interference. On the other hand, a small contention probability would restrict the number of links participate in the transmission. In summary, PHY-aware distributed scheduling with MPR requires that each link makes decision to transmit or not individually under the existence of multiple random factors, namely the SINR condition, the number of contending links, and the number of mini-slots required for channel probing, and hence is challenging.

Let S denote the set consisting of contending links. Under the MPR model, when one or more data packets are transmitted, each of them has certain probability of being received successfully, depending on the channel condition and the strength of the cochannel interference. Particularly, let $s_i = 1$ indicate that link i transmits a data packet, while $s_i = 0$ indicates no data transmission from link i. Let $r_i = 1$ indicate a successful reception of the data packet for link i, and $r_i = 0$ otherwise. Let K be a set of links such that $K \subseteq S$. The conditional probability of successful reception $q_{S,K}$ is defined as

$$q_{S,K} = Pr(r_{i\in K} = 1, r_{i\notin K} = 0 | s_{i\in S} = 1, s_{i\notin S} = 0). \quad (11)$$

Let P denote the transmission power, G_{ij} denote the channel gain from the ith transmitter to the jth receiver. The signal-to-interference-plus-noise ratio (SINR) for link i is given by $\mathrm{SINR_i} = G_{ii}P / \sum_{j\in S/i} G_{ji}P + \eta_i$, where η_i is the power spectrum density of thermal noise. It is clear that the transmission rate is an increasing function of SINR. In practical systems, the rates are often quantized to discrete values. In this section, we let $\{R^l, l = 1, 2, \cdots, L\}$ denote the set of possible discrete rates.

Under the MPR model, multiple links can transmit successfully simultaneously, and the total rate would be the sum of the rates of all successful links that participate in the transmission. We have the following result on the optimal PHY-aware scheduling policy that maximizes the average throughput.

Proposition 3.1: (**Physical interference model**) The optimal stopping rule N^* for PHY-aware distributed scheduling exists, and is given by

$$N^* = \min\left\{n \geq 1 : \max_i\{R_{i,n}\} \geq \alpha^* - E[Y_n]\right\}. \quad (12)$$

where $Y_n = \sum_{i\in K_n} R_{i,n} - \max_i\{R_{i,n}\}$, and α^* is the unique solution to

$$E(\sum_{i\in K_n} R_{i,n} - \alpha)^+ = \frac{\alpha\tau}{T}. \quad (13)$$

Remarks: Proposition 3.1 reveals that the optimal stopping rule for PHY-aware distributed scheduling under the MPR model is a pure threshold policy, and different links have different rate thresholds in general. This is in stark contrast to the scheduling for the collision model, where the optimal thresholds are the same across different links even they have different channel statistics.

IV. Channel-Aware Distributed Scheduling With Imperfect Probing

In the above, we assume that the channel state information (CSI) is perfectly known at the receiver/transmitter after channel probing. In practical scenarios, channel conditions are often estimated using noisy observations. Therefore, it is of great interest to study channel-aware scheduling under noisy channel estimation.

A. Threshold Structure with Imperfect Channel Estimation

In this section, we establish the optimal scheduling for the noisy channel estimation case. It is clear that the actual SNR λ_n is no greater than the estimated SNR $\rho_{eff}\hat{\lambda}_n$. As a result, if the packet is transmitted at the estimated rate $\log(1 + \rho_{eff}\hat{\lambda}_n)$, a channel outage will occur. Therefore, the transmission rate has to back off from the estimated rate. Equivalently, we can back off the estimated SNR $\rho_{eff}\hat{\lambda}$ to a "nominated" SNR $\lambda_c(\hat{\lambda})$, where $\lambda_c(\cdot)$ is a backoff rate function satisfying the following inequality: $0 \leq \lambda_c(\hat{\lambda}_n) \leq \lambda_n$.

Accordingly, the instantaneous rate with backoff, $R_n^{(\mathrm{BK})}$, is given by

$$R_n^{(\mathrm{BK})} = \log\left(1+\lambda_c(\hat{\lambda}_n)\right)\mathbf{I}\left(\lambda_c(\hat{\lambda}_n) \leq \lambda_n\right). \quad (14)$$

We note that there are two major differences between the perfect and noisy channel estimation cases. First, the stopping rule N in the noisy channel estimation case is defined over the filtration $\{\mathcal{F}'_n\}$ (instead of $\{\mathcal{F}_n\}$), generated by $\{(\rho|\hat{h}_j|^2, K_j), j = 1, 2, \ldots, n\}$. Second, the instantaneous rate with backoff, $R_n^{(\mathrm{BK})}$, defined in (14), is now a r.v., and is not perfectly known at time n. However, it can be easily shown that the structure of the optimal scheduling strategy remains the same, except that the random "reward" $R_n^{(\mathrm{BK})}$ is replaced with its conditional expectation, $\bar{R}_n^{(\mathrm{BK})} \triangleq E\left[R_n^{(\mathrm{BK})}|\mathcal{F}'_n\right]$ [15][18]. More specifically, we define

$$Q' \triangleq \{N \geq 1 : \{N = n\} \in \mathcal{F}'_n, E[T_N] < \infty\}, \quad (15)$$

$$Q'' \triangleq \{N \geq 1 : \{N = n\} \in \mathcal{F}''_n, E[T_N] < \infty\}, \quad (16)$$

where $\mathcal{F}''_n$ is the σ-field generated by $\{(\bar{R}_j^{(\mathrm{BK})}, K_j), j = 1, 2, \ldots, n\}$. We have the following proposition[19].

Proposition 4.1: (**Noisy Probing model**)

$$\sup_{N\in Q'} \frac{E[R_N^{(\mathrm{BK})}T]}{E[T_N]} = \sup_{N\in Q''} \frac{E[\bar{R}_N^{(\mathrm{BK})}T]}{E[T_N]}. \quad (17)$$

Proposition 4.1 indicates that the optimal scheduling can be based solely on $\bar{R}_n^{(\mathrm{BK})}$, the conditional expectation of $R_n^{(\mathrm{BK})}$ given $\mathcal{F}'_n$. In summary, we can conclude that the optimal scheduling policy under noisy channel estimation is still a pure threshold policy.

B. Threshold Structure with Two-Levels of Imperfect Channel Estimation

Consider further the case with imperfect channel information. Since the channel information is noisy, a natural question is whether it is worthwhile carrying our further channel probing to improve the quality of channel estimation. This is of particular interest in the wideband regime where the SNR is very small (i.e., $\rho = \Theta(1/W)$) and the rate is of $\Theta(1)$. Specifically, after the n-th successful channel probing, the corresponding user obtains first level information $\hat{h}_n$ of channel condition h_n. Then, it has one of the following options:

1) Accept the current estimate of the channel state and transmit at rate $\bar{R}_n$.
2) Gives up the current opportunity and let all the nodes re-contend.
3) Given the current channel information, carry out further (2nd-level) channel estimation, via sending another pilot packet, and acquire finer estimate $\hat{h}'_n$ of the channel condition. It can then decide to transmit or to defer based on $\bar{R}'_n$.

Let r_0 denoted the expected throughput. Based on [20], we can show that the optimal strategy has a two-threshold structure.

Proposition 4.2: (**Two-level noisy probing model**) After each successful channel contention, one of the following two statements holds:

1) There exist two positive constants $\bar{R}_l \leq \bar{R}_u$, such that it is optimal i) to transmit immediately after the first level of channel estimation if $\bar{R}_n > \bar{R}_u$; ii) or to give up the transmission and let all the nodes re-contend if $\bar{R}_n < \bar{R}_l$; or iii) to engage in the 2nd-level probing if $\bar{R}_n \in \left[\bar{R}_l, \bar{R}_u\right]$, and then to transmit at the rate $\bar{R}'_n$ if $\bar{R}'_n > r_0$ and to give up the transmission if $\bar{R}'_n < r_0$.
2) It is never optimal to use 2nd-level channel estimation. It is optimal to transmit at a rate $\bar{R}_n$ immediately after the first level probing if $\bar{R}_n > r_0$ and to defer transmission and re-contend if $\bar{R}_n < r_0$.

V. Channel-Aware Distributed Scheduling in Multi-hop Wireless Networks: Threshold Policies and Delay Performance

So far, our studies have focused on single-hop ad-hoc networks. Needless to say, the development of channel-aware distributed scheduling is much more challenging in multi-hop wireless networks. Note that many distributed scheduling algorithms have been devised for multi-hop networks since the seminal work [3], including max-weight scheduling, pick-and-compare scheduling, and maximal/greedy scheduling – to name a few. The complexity level of these wireless scheduling algorithms ranges from constant-time to exponential, and therefore requires a lot of message passing between communication links/nodes, hindering them from being implemented. It is highly desirable to develop threshold-based distributed scheduling in multi-hop wireless networks.

Suppose that all nodes in the network have infinite buffers, and can store a sequence of (possibly infinite many) pre-defined thresholds, which be determined offline based the the channel statistics and queueing lengths. Motivated by the max-weight scheduling, one can carry out quantization of the product of the queue length and the channel rate, μQ, over its full range, and form a sequence of thresholds $\{x_k, k = 1, 2, \ldots\}$. (This is in analogous to universal quantization.) We have the following approximation algorithm for max-weight scheduling: For each link l, whenever its local weight μQ crosses (either up or down) a threshold x_k, it notifies its neighboring links; and the "latest" max-scheduling is re-computed. One main difference between this algorithm and max-weight scheduling is that the search for max-weight matching is triggered only when the thresholds are crossed and the scheduling is updated less frequently, indicating that the complexity is much lower. It is not difficult to show that this algorithm is asymptotically throughput-optimal as long as the maximum difference between two thresholds is bounded. it remains open, however, to characterize the optimal quantization of μQ.

We note that most studies on distributed scheduling are concerned with throughput. By and large, the delay performance corresponding to wireless scheduling is an under-explored area, partially due to the fact that it is very challenging. In particular, the service rates depend on the queue sizes, which in turn depend on the arrival rates and channel conditions; and it is this coupling, between arrival rate and service rate, that complicates the characterization of delay performance. Needless to say, precise modeling of large-scale networks is prohibitive, and a sensible simplification is the key to obtain a conceptually clear understanding of the delay performance.

It is clear that standard queueing analysis does not work well for large-scale distributed networks, because it is impossible to characterize the service distribution due to the coupling between the queue size and the scheduling. Among other tools for delay analysis, Lyapunov function theory can provide bounds only in general, and the large-deviation approach is often limited by the dimensionality of the problem. In light of this, we are currently developing heavy traffic theory to analyze the delay performance.

VI. Acknowledgements

The author would like to thank fruitful collaborations with Dong Zheng, Weiyan Ge, Chandra P.S.B, P. R. Kumar, Simon Pun, H. V. Poor, Yung Yi, Mung Chiang, and helpful discussions with Ness Shroff, R. Srikant, Armand Makowski and Anthony Ephremides.

References

[1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *Proceedings of SIGCOMM*, 2004.

[2] M. Cao, V. Raghunathan, and P. R. Kumar, "Cross layer exploitation of MAC layer diversity in wireless networks," in *Proceedings of IEEE ICNP*, 2006.

[3] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Transactions on Information Theory*, vol. 39, pp. 466–478, March 1993.

[4] G. Sharma, R. R. Mazumdar, and N. B. Shroff, "On the complexity of scheduling in wireless networks," in *Proceedings of MOBICOM*, 2006.

[5] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, P. Whiting, and R. Vijayakumar, "Providing quality of service over a shared wireless link," *IEEE Comm. Magazine*, vol. 39, pp. 150–154, 2001.

[6] S. Borst, "User-level performance of channel-aware scheduling algotirthms in wireless data networks," in *Proc. IEEE INFOCOM'03*, 2003.

[7] R. Knopp and P. Humlet, "Information capacity and power control in single cell multiuser communications," in *Proc. IEEE ICC 95*, June 1995.

[8] X. Liu, E. K. Chong, and N. B. Shroff, "A framework for opportunistic scheduling in wireless networks," *Computer Networks*, vol. 41, no. 4, pp. 451–474, Mar. 2003.

[9] S. Shakkottai, R. Srikant, and A. L. Stolyar, "Pathwise optimality and state space collapse for the exponential rule," in *Proceedings of IEEE Symposium on Information Theory*, July 2002.

[10] D. Tse, "Multiuser diversity in wireless networks," http://degas.eecs.berkeley.edu/~dtse/pub.html, Apr. 2001.

[11] P. Viswanath, D. N. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Trans. Info. Theory*, vol. 48, no. 6, pp. 1277–1294, June 2002.

[12] X. Qin and R. Berry, "Exploiting multiuser dieversity for medium access control in wireless networks," in *Proceedings of IEEE INFOCOM'03*, 2003.

[13] S. Adireddy and L. Tong, "Exploiting decentralized channel state information for random access," *IEEE Trans. Info. Theory*, vol. 51, no. 2, pp. 537–561, Feb. 2005.

[14] Y. Chow, H. Robbins, and D. Siegmund, *Great Expectations: Theory of Optimal Stopping*. Houghton Mifflin, 1971.

[15] T. Ferguson, *Optimal Stopping and Applications*. Online.

[16] D. Zheng, W. Ge, and J. Zhang, "Distributed opportunistic scheduling for ad-hoc networks with random access: An optimal stopping approach," *IEEE Transactions on Information Theory*, submitted, 2007.

[17] D. Zheng, M. Cao, J. Zhang, and P. R. Kumar, "Channel aware distributed scheduling for exploiting multiuser diversity and multi-receiver diversity in ad-hoc networks: A unified PHY/MAC approach," in *Proc. IEEE INFOCOM'08*, 2008.

[18] D. Assal, L. Coldstein, and E. Samuel-Cahn, "A statistical version of prophet inequalities," *The Annals of Statistics*, vol. 26, no. 3, pp. 1190–1197, 1998.

[19] D. Zheng, S. Pun, W. Ge, J. Zhang, and H. Poor, "Distributed opportunistic scheduling for ad-hoc communications under noisy channel estimation," in *Proceedings of ICC*, 2008.

[20] W. Stadje, "An optimal stopping problem with two levels of incomplete information," *Mathematical Methods of Operations Research*, vol. 45, pp. 119–131, 1997.