

Optimal Opportunistic Scheduling in Wireless Networks

Xin Liu

Department of Computer Science
University of California, Davis, CA 95616-8562

Edwin K. P. Chong

Dept. of Electrical and Computer Engineering
Colorado State University, Fort Collins, CO 80523

Email: echong@engr.colostate.edu

Ness B. Shroff

School of Electrical and Computer Engineering
Purdue University, West Lafayette, IN 47907-1285

Email: shroff@ecn.purdue.edu

Abstract—Scheduling has been extensively studied in various disciplines in operations research and wireline networking. However, the unique characteristics of wireless communication systems – namely, timing-varying channel conditions and multi-user diversity – means that new scheduling solutions need to be developed that are specifically tailored for this environment. In this paper, we summarize various opportunistic scheduling schemes that exploit the time-varying nature of the radio environment to improve the spectrum efficiency while maintaining a certain level of satisfaction for each user. We also discuss the advantages and costs associated with opportunistic scheduling, and identify possible future research directions.

I. INTRODUCTION

Transmission scheduling for wireless networks has recently attracted a lot of research interests. First, scheduling policies of wireline networks are extended to wireless networks, where the burst of errors in wireless channels is taken into account. To elaborate, a wireless channel is modeled by a two-state Markov chain [1]: a user experiences error-free transmission when it observes a “good” channel, and unsuccessful transmission in a “bad” channel. Using such a channel model, wireless fair scheduling policies have been studied [2], [3], [4], [5]. These works provide various degrees of performance guarantees, including short-term and long-term fairness, as well as short-term and long-term throughput bounds. A survey of these algorithms can be found in [6]. The limitation of these works is that channels are modeled as either “good” or “bad,” which is too simple to characterize realistic wireless channels, especially for data services.

The IS-856 system, which is also known as High Data Rate (HDR) [7], has been developed at Qualcomm to provide a versatile wireless Internet solution [8]. The scheduler in IS-856 is proportional fair [9], [10]. To elaborate, the proportional fairness (PF) scheduler maximizes the product of the

throughput delivered to all the users. In other words, the set of throughput achieved by different users is proportionally fair if increasing the throughput of one user from the current level by $x\%$ requires a cumulative percentage decrease in all the users of more than $x\%$. Suppose that there are N users and $C_i(t)$ is the estimate of the average rate for user i at time-slot t , $i = 1, \dots, N$. Also, suppose that at time-slot t , the current achievable data rate of user i is $R_i(t)$, $i = 1, \dots, N$. The algorithm works as follows:

- Scheduling: The user with the highest ratio of $R_i(t)/C_i(t)$ of all N users will receive transmission at each decision time. Ties are broken randomly.
- Update average rate: For each user i ,

$$C_i(t+1) = (1 - 1/t_c)C_i(t) + 1/t_c \times R_i(t) \times \mathbf{1}_i,$$

where $\mathbf{1}_i = 1$ if user i is chosen to transmit, otherwise $\mathbf{1}_i = 0$.

The value of parameter t_c used by the scheduling algorithm is related to the maximum amount of time for which an individual user can be starved [11]. The PF scheduling scheme is analyzed in [12], [13].

In [14], [15], [16], the authors study scheduling algorithms for the transmission of data to multiple users. Both delay and channel conditions are taken into account. Roughly speaking, the algorithm can be described as $\operatorname{argmax} \rho_i W_i R_i$, where W_i is the head-of-the-line packet delay for queue i , R_i is the data rate of user i , and ρ_i is some constant. The proposed scheduler achieves throughput optimality, defined in [14] as follows: a scheduling algorithm is throughput optimal if it is able to keep all queues stable if this is at all feasible to do with any scheduling algorithm. Further, in [17], [18], the authors propose a scheduling scheme with an exponential rule that well balances the tradeoff between queue length and throughput. The exponential-rule scheduler is also throughput optimal.

The authors of [19] investigate a scheduling algorithm to maximize the minimum (weighted) throughput of users. The

This research is supported in part by NSF awards ANI-0207728, ANI-0099137, EIA-0130599, ECS-0098089 and ANI-0207892, and the Indiana 21st century center for wireless communications and networking.

optimal solution is in the form of $\operatorname{argmax}_i c_i R_i(t)$, where c_i can be interpreted as the shadow price or reward, whose value depends on the distributions of R_i . The authors also propose an adaptive algorithm to determine the parameters, and study the transient behavior.

Opportunistic scheduling exploits the channel fluctuations of users. Hence, the larger the channel fluctuation, the higher the scheduling gain. Thus a natural question to ask is what we should do in environments with little scattering and/or slow fading. In [20], the authors use multiple transmission antennas to “induce” channel fluctuations, and thus exploit multi-user diversity in a slow fading environment. Specifically, the antennas are fed with randomly picked phase and amplitude to “induce” channel fluctuation. Each user feeds back the overall SINR of its “induced” channel to the base station. The base station selects the user with a large peak value of SINR to transmit according to a certain scheduling rule. When there are a large number of users, the base station can always find a user with its peak SINR to transmit. Hence, the system performance is asymptotically as good as a solution with an optimal beam-forming configuration, while using only the overall SINR as feedback. Further, such a scheme can also be used opportunistically to null intercell interference.

In [21], the author studies the opportunistic scheduling from a flow-level. Instead of assuming infinite backlogs, the author studies the dynamic case where users have random finite-size service demands. The authors show that under certain assumptions, the user-level performance can be approximated by a multi-class Processor-sharing model where the total service rate varies with the total number of users.

In [22], power consumption of users in a fading channel is considered. In general, by varying the transmission rate and power, based on the current fading level, a user in the wireless network can utilize the available energy more efficiently. However, such an approach can lead to long delays or buffer overflows. In [22], the tradeoffs between the required power and various notions of delay are analyzed.

In [23], the authors study scheduling problems for real-time traffic with fixed deadlines. The authors propose a greedy algorithm that chooses the request with the largest revenue in the current time-slot to serve, and show that the greedy algorithm is 1/2 competitive against the offline optimal algorithm; i.e., the throughput of the greedy algorithm is at least 1/2 of that of the offline optimal algorithm. Further, they show that no deterministic online algorithm can achieve a competitive ratio higher than 1/2.

Downlink scheduling in CDMA systems for data transmission is also studied in [24]. The work considers a performance metric called “stretch”, which can be considered as the normalized delay. A near optimal, offline, polynomial time algorithm is proposed to minimize the maximum stretch under the assumption of continuous rates, and various online algorithms for continuous-rates/discrete-rates are studied with simulations.

In [25], the authors study transmission schemes for time-varying wireless channels with partial state information. The objective is to minimize a discounted infinite-horizon cost function, which can be used to indicate the balance between

power cost and throughput. The resulting optimal solution is a threshold back-off scheme.

In [26], the authors present a framework for opportunistically scheduling user transmissions to exploit the time-varying channel conditions in wireless communication systems. The objective is to maximize the wireless system performance while satisfying various QoS requirements. The framework enables us to investigate different categories of scheduling problems involving two fairness requirements (*temporal fairness* and *utilitarian fairness*) and a minimum-performance requirement. In particular, temporal fairness means that each user obtains a certain portion of the resource (e.g., time-slots). Utilitarian fairness means that each user obtains a certain portion of the overall system performance. Last, a user can also have a minimum-performance requirement (e.g., a data rate of 100Kbps). The optimal scheduling solutions for these scheduling problems turn out to be index policies, and a stochastic-approximation-based algorithm can be used to efficiently estimate the key parameters of the scheduling schemes on-line. The feasibility problem of a network with minimum-performance requirements is also studied.

II. A NON-PROBABILISTIC APPROACH

In many previous problem formulations, only stationary policies are considered. (A policy is a stationary policy if it is not a function of time.) In this section, we use the temporal fairness scheduling problem as an example to show how to generalize the result to more general cases. Similar extensions hold for other scheduling problems.

Let U_i^k be the performance value of user i at time k . An example of the performance value is the data rate. Let Q be a general policy whose value at time k may depend on the entire performance value sequence $\{\vec{U}^k, k = 1, 2, \dots\}$ and the time. Let $F_i^K(Q)$ be the average performance value of user i up to time K and $T_i^K(Q)$ be the average resource consumption of user i up to time K . To elaborate,

$$F_i^K(Q) = \frac{1}{K} \sum_{k=1}^K U_i^k \mathbf{1}_{\{Q^k=i\}}, \quad i = 1, 2, \dots, N$$

$$T_i^K(Q) = \frac{1}{K} \sum_{k=1}^K \mathbf{1}_{\{Q^k=i\}}, \quad i = 1, 2, \dots, N,$$

where $Q^k = Q(\{\vec{U}^t, t = 1, 2, \dots\}, k)$, and $Q^k = i$ means that user i is scheduled to transmit at time k . The temporal fairness constraint is formulated as:

$$\liminf_{K \rightarrow \infty} T_i^K(Q) \geq t_i, \quad i = 1, 2, \dots, N,$$

where t_i is a predetermined fairness parameter, $t_i \geq 0$, and $\sum_{i=1}^N t_i \leq 1$. It says, in the worst case, at least t_i portion of the resource (time-slots) is allocated to user i asymptotically.

Let $F^K(Q) = \sum_{i=1}^N F_i^K(Q)$; i.e., $F^K(Q)$ is the average system performance up to time K . We define

$$F(Q) = \limsup_{K \rightarrow \infty} F^K(Q),$$

which can be considered as the asymptotic best-case system performance of policy Q .

Let Θ be the set of all scheduling policies, including non-stationary policies and non-causal policies. The temporal scheduling problem is formulated as:

$$\begin{aligned} & \underset{Q \in \Theta}{\text{maximize}} && F(Q) \\ & \text{subject to} && \liminf_{K \rightarrow \infty} T_i^K(Q) \geq t_i, \quad i = 1, 2, \dots, N. \end{aligned} \quad (1)$$

We state our policy Q^* as follows:

$$Q^*(\vec{U}^k) = \underset{i}{\text{argmax}} (U_i^k + v_i^*), \quad (2)$$

where the v_i^* s are chosen such that:

- 1) $\min_i (v_i^*) = 0$
- 2) $\liminf_{K \rightarrow \infty} T_i^K(Q^*) \geq t_i$ for all i
- 3) For all i , if $\liminf_{K \rightarrow \infty} T_i^K(Q^*) > t_i$, then $v_i^* = 0$.

Proposition 1: If $\lim_{K \rightarrow \infty} T_i^K(Q^*)$ exists for all i for the Q^* defined in (2), then the policy Q^* is a solution to the problem defined in (1).

Before we prove the above proposition, we explain the proposition under various scenarios.

- Suppose that $\{\vec{U}^k, k = 1, 2, \dots\}$ is stationary and ergodic. Because Q^* is a stationary policy, $T_i^K(Q^*)$ and $F_i^K(Q^*)$ converge to a constant almost surely. Thus, Q^* is a solution to the problem defined in (1). Furthermore, we have

$$\liminf_{K \rightarrow \infty} F^K(Q^*) = \limsup_{K \rightarrow \infty} F^K(Q^*). \quad (3)$$

This equation is critical. It states the important fact that the asymptotic worst-case system performance of our policy Q^* ($\liminf_{K \rightarrow \infty} F^K(Q^*)$) is the same as its asymptotic best-case system performance ($\limsup_{K \rightarrow \infty} F^K(Q^*)$). Thus, the *worst-case* performance of Q^* asymptotically bounds the *best-case* system performance of an arbitrary policy that satisfies the temporal fairness constraint.

- Suppose that $\{\vec{U}^k, k = 1, 2, \dots\}$ is stationary and ergodic. Then many policies have the nice property that $F_i^K(Q)$ and $T_i^K(Q)$ converge (to a random variable almost surely). Examples of such policies are stationary policies and periodic policies¹. However, there exist policies such that

$$\limsup_{K \rightarrow \infty} F_i^K(Q) > \liminf_{K \rightarrow \infty} F_i^K(Q).$$

In this case, even if the policy Q is a solution to the problem defined in (1), it is not a “good” solution because only its asymptotic *best-case* performance bounds that of others. On the contrary, the asymptotic *worst-case* performance of Q^* (defined in (2)) bounds the asymptotic best-case performance of others.

- The proposition holds without the assumption that $\{\vec{U}^k, k = 1, 2, \dots\}$ is stationary and ergodic. However, in this case, we may not be able to estimate the parameters (v_i^*) used in Q^* in practice.
- A round-robin policy is often used as an example of non-opportunistic policies for comparison. To be specific, round-robin is a non-stationary non-opportunistic

scheduling policy. If $\{\vec{U}^k, k = 1, 2, \dots\}$ is stationary, then the expectation of the long-term average of the performance value of a round-robin policy is equivalent to that of a non-opportunistic policy (i.e., $E(U_i)$ for user i).

Proof: For technical simplicity, we assume that v_i^* s are bounded. Furthermore, if $\sum_{i=1}^N v_i^* = 0$, then $v_i^* = 0$ for all i . In this case, Q^* always chooses the user with the maximum performance value to transmit, and thus the result is trivial. Now we consider the case where $\sum_{i=1}^N v_i^* > 0$.

If policy Q satisfies the fairness constraints; i.e., $\liminf_{K \rightarrow \infty} T_i^K(Q) \geq t_i$ for all i , then for any $\epsilon > 0$, there exists L_1 , such that for any $K > L_1$, we have

$$T_i^K(Q) > t_i - \frac{\epsilon}{2 \sum_{i=1}^N v_i^*}, \quad i = 1, 2, \dots, N. \quad (4)$$

Because of the hypothesis that $\lim_{K \rightarrow \infty} T_i^K(Q^*)$ exists and the condition 3 above to choose v_i^* s, we have

$$v_i^* \left(\lim_{K \rightarrow \infty} T_i^K(Q^*) - t_i \right) = 0, \quad i = 1, 2, \dots, N.$$

Hence, for the ϵ in (4), there exists $L > L_1$, such that for $K > L$, we have

$$|v_i^* (T_i^K(Q^*) - t_i)| < \frac{\epsilon}{2N}, \quad i = 1, 2, \dots, N.$$

Then for $K > L$, we have

$$\begin{aligned} & F^K(Q) \\ & \leq F^K(Q) + \sum_{i=1}^N v_i^* \left(T_i^K(Q) - t_i + \frac{\epsilon}{2 \sum_{i=1}^N v_i^*} \right) \\ & = \sum_{i=1}^N \frac{1}{K} \sum_{k=1}^K (U_i^k + v_i^*) \mathbf{1}_{\{Q^k=i\}} - \sum_{i=1}^N v_i^* t_i + \frac{\epsilon}{2}. \end{aligned}$$

By the definition of Q^* , we have

$$\sum_{i=1}^N (U_i^k + v_i^*) \mathbf{1}_{\{Q^k=i\}} \leq \sum_{i=1}^N (U_i^k + v_i^*) \mathbf{1}_{\{Q^*(\vec{U}^k)=i\}}.$$

Thus,

$$\begin{aligned} & F^K(Q) \\ & \leq \sum_{i=1}^N \frac{1}{K} \sum_{k=1}^K (U_i^k + v_i^*) \mathbf{1}_{\{Q^*(\vec{U}^k)=i\}} - \sum_{i=1}^N v_i^* t_i + \frac{\epsilon}{2} \\ & = F^K(Q^*) + \sum_{i=1}^N v_i^* (T_i^K(Q^*) - t_i) + \frac{\epsilon}{2} \\ & \leq F^K(Q^*) + \epsilon. \end{aligned}$$

Because ϵ is chosen arbitrarily, we have

$$\limsup_{K \rightarrow \infty} F^K(Q) \leq \limsup_{K \rightarrow \infty} F^K(Q^*),$$

which completes the proof. \square

¹This can be shown by directly applying Birkoff's ergodic theorem.

III. DISCUSSIONS

To meet the increasing demand for wireless services, especially affordable wireless Internet services, wireless spectrum efficiency is becoming increasingly important. In wireless networks, users experience unreliable, location-dependent, and time-varying channel conditions. Traditionally, the channel variation is considered as a negative factor for reliable communication, and should be mitigated by methods such as time interleaving, power control, and multiple antennas. On the other hand, opportunistic scheduling is designed to exploit the variation of channel conditions to improve spectrum efficiency. It adds an additional degree of freedom to the system: time-domain diversity or also called multiuser diversity. It improves spectrum efficiency, especially for delay-tolerant data transmissions. Various opportunistic scheduling schemes have been studied. A common objective is to improve/maximize system performance (e.g., throughput) under various fairness and QoS constraints. In many cases, the optimal policies are given in a simple parametric form, hence lending themselves to easy implementations. The advantages of opportunistic scheduling also include the ability to work with other resource management mechanisms. A good example of this is the joint scheduling and power-allocation scheme [27].

However, nothing comes for free. Opportunistic scheduling also has its own costs and limitations. Many

- There are signaling costs involved in all opportunistic scheduling schemes because scheduling decisions inherently depend on channel conditions (and/or queueing status). Users need to constantly estimate their channel conditions and report to the base station. Hence, the actual scheduling gain should take into account the signaling costs.
- Because users need to estimate the channel conditions, estimation errors occur in all scheduling schemes. There are various sources of estimation errors: errors of estimations of channels, errors of estimations of parameters involved in scheduling schemes, and errors caused by various delays such as transmission delay, estimation delay, and restriction of time-slots, etc. In general, if the variation of channel conditions is relatively slow, then the estimation is good. We recommend a rigorous study on this problem, especially in the case of fast fading.
- Opportunistic scheduling exploits the fluctuation of channel conditions, and thus scheduling gain inherently depends on the amplitude of the variations of channels. In general, the greater the fluctuation of channel conditions, the larger the number of users, the better the performance gain.
- Another concern in opportunistic scheduling is the time scale of fluctuation. The fluctuation of channels should be slow enough for user to estimate it and exploit it. On the other hand, the fluctuation should be fast enough, so that users won't experience extreme long delays. (Though many data users are delay-tolerant, extreme delays may cause upper-layer problems such as TCP timeout.)
- There is a tradeoff between scheduling gain and short-term performance. In general, the stronger the time-

correlation of channel conditions (i.e., the slower the channel fluctuation), the worse the short-term performance, and the greater the improvement in the short-term performance, the less the scheduling gain.

- It is reported that the scheduling gain may decrease when there are multiple antennas. Because smart antennas, including MIMO, are promising technologies in future generation wireless networks, the relationship between opportunistic scheduling and antennas arrays should be further studied.

In summary, opportunistic scheduling presents a new design approach, especially for delay-tolerant data traffic. It has its own advantages and limitations. Many interesting problems are yet to be resolved in this area. We discuss some possible research problems in the next section.

IV. POSSIBLE RESEARCH DIRECTIONS

Various long-term fairness criteria, such as proportional fairness, temporal fairness, and utilitarian fairness, have been studied for scheduling problems in wireless networks. However, there is a need for general short-term fairness criteria tailored to wireless networks and dealing with the short-term performance in depth. References related to the subject include [14], [15], [16], [17], [18], where queueing status is a part of scheduling decisions.

A problem related to improving short-term performance is to schedule traffic with deadlines, i.e., real-time traffic. Specifically, upon arrival, each real-time packet has a delay deadline, and packets that cannot be transmitted before their deadlines are dropped/marked. Research on scheduling with deadlines in the wireline setting has led to various approaches. The goal is typically to minimize some measure of the number of deadline misses (including weighting such misses according to packet classes, also called *weighted loss*). The additional challenge in wireless networks is due to the time-varying channel conditions. In this type of problems, the objective is to improve system performance (with or without fairness/QoS constraints) by exploiting multi-user diversity. Approaches to these problems may include off-line optimal solutions with the assumption of entire traffic and channel information, on-line model-based solutions, and heuristic/greedy algorithms [23]. Heuristic algorithms play an important role in real-time scheduling problems because (typically) the optimal scheduling problem is NP-complete and simplicity is a desirable feature. In the wireline world, it is sometimes the case that complicated scheduling schemes do not have significant performance gains over simple schemes, such as static priority or earliest-deadline-first. A similar situation may be expected to hold for wireless networks.

The opportunistic scheduling is based on the premise that the wireless channel is time-varying, and we can schedule users to transmit at those times that are opportunistically "relatively good." This idea can be extended to the frequency domain: we opportunistically schedule users to frequencies (and time) that are relatively good [28]. An example of such systems is an OFDM system. A concern of opportunistic scheduling in such systems is the signaling cost. Because each

subcarrier is very narrow in OFDM systems, signaling should be carefully designed to ensure good channel estimation of users on different subcarriers while avoiding significant signaling overhead.

The opportunistic scheduling scheme in its current form is a network-layer problem. However, its performance is closely related to physical-layer designs. As explained earlier, estimation errors occur in all opportunistic scheduling schemes. On one hand, we need better understandings of the effect of channel estimation errors on scheduling schemes. On the other hand, it calls for better channel estimation techniques and smart coding schemes (e.g., incremental redundancy transmission schemes with turbo codes). Further, it is also important to study the performance of opportunistic scheduling in multiple antenna systems. In summary, a better understanding of physical-layer technologies or even layer-breaking designs can be potentially beneficial.

In general, the scheduling gain increases as the number of users increases. However, the normalized scheduling gain (scheduling gain over number of users) decreases with the increase of the number of users. For example, if U_i 's are i.i.d. with exponential distribution, then the scheduling gain is $O(\log(n))$ [20]. On the other hand, the signaling cost per user remains the same. Hence, it is a question of practical importance to decide the number of users sharing a same channel.

The opportunistic scheduling problems studied have the net effect of increasing the overall effective capacity of the wireless network. This means that the network can now accommodate more users or higher-data-rate users. Thus, we know that keeping all else constant, the *admissible region* of the wireless network will increase by using opportunistic scheduling schemes. A challenging problem that still remains is making intelligent *admission control* decisions of whether or not to allow a new user into a cell. Although admission control is a difficult problem in wireless systems whether or not opportunistic scheduling is used, it is more challenging in the context of opportunistic scheduling because opportunistic scheduling increases the system dynamics.

Most of the current research on opportunistic scheduling focus on the downlink of a cellular system. In such a system, there exists a natural central controller, the base station. An interesting question is whether and how to exploit the time-domain diversity in a distributed environment, such as an ad-hoc network [29]. Another challenging question is how to implement opportunistic scheduling in the upper-link of a cellular system.

REFERENCES

- [1] E. Gilbert, "Capacity of a burst-noise channel," *Bell Systems Technical Journal*, vol. 39, pp. 1253–1266, 1960.
- [2] S. Lu, V. Bharghavan, and R. Srikant, "Fair scheduling in wireless packet networks," *IEEE/ACM Transactions on Networking*, vol. 7, no. 4, pp. 473–489, August 1999.
- [3] T. Nandagopal, S. Lu, and V. Bharghavan, "A unified architecture for the design and evaluation of wireless fair queueing algorithms," in *Proceedings of ACM Mobicom 1999*, August 1999.
- [4] T. Ng, I. Stoica, and H. Zhang, "Packet fair queueing algorithms for wireless networks with location-dependent errors," in *Proceedings of IEEE INFOCOM 1998*, New York, 1998, vol. 3.

- [5] S. Shakkottai and R. Srikant, "Scheduling real-time traffic with deadlines over a wireless channel," in *Proceedings of ACM Workshop on Wireless and Mobile Multimedia*, Seattle, WA, August 1999.
- [6] Y. Cao and V. Li, "Scheduling algorithms in broadband wireless networks," *Proceedings of the IEEE*, vol. 89, no. 1, pp. 76–87, January 2001.
- [7] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushyana, and A. Viterbi, "CDMA/HDR: a bandwidth-efficient high-speed wireless data service for nomadic users," *IEEE Communications Magazine*, vol. 38, no. 7, pp. 70–77, July 2000.
- [8] "1xEV: 1x EVolution, IS-856 TIA/EIA standard," *QUALCOMM Inc.*
- [9] A. Jalali, R. Padovani, and R. Pankaj, "Data throughput of CDMA-HDR a high efficiency-high data rate personal communication wireless system," in *Proceedings of IEEE Vehicular Technology Conference 2000-Spring*, 2000, vol. 3.
- [10] F. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, vol. 8, pp. 33–37, 1997.
- [11] "1x high data rate airlink overview," *QUALCOMM Inc.*
- [12] J. Holtzman, "CDMA forward link waterfilling power control," in *Proceedings of IEEE Vehicular Technology Conference 2000-Spring*, 2000, vol. 3, IEEE.
- [13] J. Holtzman, "Asymptotic analysis of proportional fair algorithm," in *Proceedings of 12th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, 2001, vol. 2.
- [14] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, P. Whiting, and R. Vijayakumar, "Providing quality of service over a shared wireless link," *IEEE Communications Magazine*, vol. 39, no. 2, pp. 150–153, February 2001.
- [15] M. Andrews, S. Borst, F. Dominique, P. Jelenkovic, K. Kumaran, K. Ramakrishnan, and P. Whiting, "Dynamic bandwidth allocation algorithms for high-speed data wireless networks," *Bell Labs Technical Report*, 2000.
- [16] S. Shakkottai and A. Stolyar, "Scheduling algorithms for a mixture of real-time and non-real-time data in HDR," *Bell Laboratories Technical Report*, 2000.
- [17] S. Shakkottai and A. Stolyar, "Scheduling for multiple flows sharing a time-varying channel: The exponential rule," *Translations of the AMS*, 2001, A volume in memory of F. Karpelevich.
- [18] S. Shakkottai and A. Stolyar, "Scheduling of a shared a time-varying channel: The exponential rule stability," in *INFORMS Applied Probability Conference*, New York, July 2001.
- [19] S. Borst and P. Whiting, "Dynamic rate control algorithms for CDMA throughput optimization," in *Proceedings of IEEE Infocom 2001*, Alaska, April 2001.
- [20] P. Viswanath, D. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Transactions on Information Theory*, vol. 48, no. 6, pp. 1277–1294, 2002.
- [21] S. Borst, "User-level performance of channel-aware scheduling algorithms in wireless data networks," in *Proceedings of IEEE INFOCOM*, 2003.
- [22] R. Berry, *Power and delay trade-offs in fading channels*, Ph.D. thesis, MIT, 2000.
- [23] M. Agarwal and A. Puri, "Base station scheduling of requests with fixed deadlines," in *Proceedings of IEEE INFOCOM 2002*, 2002, IEEE.
- [24] N. Joshi, S. Kadaba, S. Patel, and G. Sundaram, "Downlink scheduling in CDMA data networks," in *Proceedings of ACM Mobicom 2000*, 2000.
- [25] D. Zhang and K. Wasserman, "Transmission schemes for time-varying wireless channels with partial state observations," in *Proceedings of IEEE INFOCOM 2002*, 2002, IEEE.
- [26] X. Liu, E. Chong, and N. Shroff, "A framework for opportunistic scheduling in wireless networks," *Computer Networks*, vol. 41, no. 4, pp. 451–474, March 2003.
- [27] X. Liu, E. Chong, and N. Shroff, "Joint scheduling and power-allocation for interference management in wireless networks," in *Proceedings of IEEE Vehicular Technology Conference 2002-Fall*, Vancouver, Canada, 2002, vol. 3.
- [28] Y. Liu and E. Knightly, "Opportunistic fair scheduling over multiple wireless channels," in *Proceedings of IEEE INFOCOM 2003*, 2003.
- [29] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," in *Proceedings of ACM MOBICOM*, 2002.