OPPORTUNISTIC SCHEDULING IN WIRELESS COMMUNICATION NETWORKS

A Thesis

Submitted to the Faculty

of

Purdue University

by

Xin Liu

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

December 2002

To my parents, Huan, and Rucheng

ACKNOWLEDGMENTS

I would like to show my sincerest gratitude to my co-advisors, Professor Edwin K. P. Chong and Professor Ness B. Shroff for giving me copious amounts of insightful guidance, constant encouragement, constructive criticism, and expertise on every subject that arose throughout all these years. Their enthusiasm and dedication to their students are truly inspiring; it is my very privilege to have been one of them. I would also like to thank Professors Kihong Park and Saul Gelfand for being on my advisory committee and for all their helps along the way.

I am very thankful for the many friends and officemates I have had at Purdue. Special thanks to Uday R. Savagaonkar, Junshan Zhang, Suresh Kalyanasundaram, Jeffrey Herdtner, Mingbo Xiao, Dongyu Qiu, Do Young Eun, Han S. Kim, Xiaojun Lin, Jeongjoon Lee, Jungmin Park, Dong-Won Shin, Jang Won Lee, Gang Wu, Qing Su, Yu Ying for many stimulating discussions and warm friendship. I would also like to thank many other friends, especially Limin Qian, Fan Wang, Scott and Jean Harford, and our volleyball team; my years at Purdue would not have been the same without them.

Finally, I want to thank my husband, Huan Ren, for his love, support, encouragement, sense of humor, and for being an intelligent colleague as well. I could not have accomplished this without him. Very special thanks to my parents for their selfless love and support, and to Eric Rucheng Ren for being my other accomplishment at Purdue.

TABLE OF CONTENTS

Page

LIST OF TABLES

LIST OF FIGURES

Figure Page

ABSTRACT

Liu, Xin. Ph.D., Purdue University, December, 2002. Opportunistic Scheduling in Wireless Communication Networks. Major Professors: Edwin K. P. Chong and Ness B. Shroff.

Wireless spectrum efficiency is becoming increasingly important with the growing demand for wideband wireless services. In this dissertation, we present a unified framework for opportunistic scheduling, which exploits the time-varying nature of the radio environment to increase the overall performance of the system. Our framework enables us to investigate various categories of scheduling problems, differing through their QoS objectives. The QoS objectives studied in this dissertation include two fairness requirements (temporal fairness and utilitarian fairness) and a minimum-performance requirement. We find optimal solutions for these scheduling problems. An attractive feature of these optimal solutions is that they are given in a simple parametric form, hence lending themselves to on-line implementation. We provide algorithms to estimate the parameters in these optimal solutions and describe implementation procedures for each solution.

An advantage of opportunistic scheduling is that it can be coupled with other resource management mechanisms to further increase network performance. In this dissertation, we study joint scheduling and power control schemes for intercell-interference alleviation. We investigate two different problems in this context. In the first problem, the objective is to minimize the average transmission power, and thus interference to other cells, while maintaining the required data-rate for each user within the cell. In the second problem, the objective is to maximize the net utility (defined as the difference between the value of throughput and the cost of power consumption) with the same data-rate requirements. We establish the optimality of our

joint scheduling and power-allocation schemes for these problems, and discuss their properties.

1. INTRODUCTION

1.1 Overview of Cellular Communication Systems

The last decade has witnessed a tremendous growth in the wireless market. First generation (analog voice) and second generation (digital voice/low-rate data) wireless networks have been ubiquitously deployed. While first generation and second generation wireless networks focus on voice services, future generation wireless networks face new challenges — high-rate-data services and QoS (Quality of Service) support. The global demand for wireless "bandwidth" exhibits, now and in the foreseeable future, significant growth [1]. Compared with wireline networks, wireless resource is very scarce. While more wired network "bandwidth" is created when new physical resources (cable, fiber, router, etc.) are added to the network, wireless communication requires sharing a finite natural resource: the radio frequency spectrum. The data-rate capacity that a radio frequency channel can support is limited by Shannon's capacity laws [2]. Significant efforts have been made to improve the wireless spectrum efficiency in order to meet the future demand for high-data-rate wireless communication. We next give a brief introduction of some key components.

Multiuser Detection

Wireless multi-user systems are subject to co-channel interference. However, multi-access noise has considerable structure, and certainly much less randomness than white Gaussian background noise. By exploiting that structure, multi-user detection can increase spectral efficiency, receiver sensitivity, and the number of potential users. The key element used in optimum multiuser detectors is a bank of matched filters (the outputs are sufficient statistics for demodulation), followed by a dynamic programming algorithm (Viterbi's algorithm). Although optimum multiuser detectors can achieve good performance, the corresponding computational complexity increases exponentially in the number of users; i.e., the optimum multiuser detection problem is NP-hard. On the other hand, linear multiuser receivers, such as decorrelators and linear minimum-mean-square-error (MMSE) receivers, are suboptimal but practically more appealing [3]. Note that the conventional matched filter receiver can be considered as a degenerate special case of a linear multiuser receiver.

Antenna arrays

Directive antennas can be used to concentrate energy in the direction of the receiver/transmitter. They have lower power requirements and minimize interference to and from other antennas. Furthermore, antenna arrays can greatly improve spectral efficiency even in environments with significant local scattering. Smart antennas (in the conventional terminology) exploit spatial diversity by optimally combining the response from each antenna element [4, 5]. Smart antennas use multiple antenna elements on one end only: SIMO (single input multiple output) on the receiver side, and MISO (multiple input single output) on the transmitter side. Further, if receiving/transmitting array elements are sufficiently separated, the fading parameters at different elements become weakly dependent. This transmitter/receiver diversity mechanism effectively creates a plurality of subchannels sharing the same RF bandwidth. MIMO (multiple input multiple output) is built on such assumptions. A good example is the Bell Laboratories BLAST system [6]. The BLAST system has demonstrated spectral efficiencies on the order of 40b/s/Hz with eight elements at both transmitter and receiver. This is more than 40 times the achievable data rate with single-element transmitting/receiving antennas using the same bandwidth and power.

Channel error control coding

Error control coding arose from the seminal contribution in communication theory made by Shannon [2] that establishes fundamental limits on reliable communication, and presents the challenge of finding specific families of codes that achieve the capacity limit. Error control codes add redundancy to the data in order to protect it from the random disturbances introduced by the channel. In the last several years, significant work has been done on Turbo coding/decoding [7]. Turbo coding can potentially achieve performance that is close to the Shannon capacity limits at the expense of complexity.

Power control

In wireless communication systems, the received power represents signal strength to the desired receiver but interference to all other users. Power control is intended to provide each user an acceptable connection by eliminating unnecessary interference. The elegant work of Yates [8] abstracts the important properties of various power control algorithms and presents a unified treatment of power control. While power control is widely implemented in CDMA systems, such as IS-95 [9], it has also been shown to increase the call carrying capacity for channelized systems, such as TDMA/FDMA systems [10]. Furthermore, beyond the conventional concept of power control as a means to eliminating the "near-far" effect, power control is an effective resource management mechanism. It plays an important role in interference management, channel-quality/service-quality provisioning, and capacity management [11, 10, 12, 13].

Admission control

Empirical studies have shown that a typical user is far more irritated when an ongoing call is dropped than a call blocked from the very beginning. Hence, a goal of admission control is to admit as many users as possible (to maximize the revenue of the system) while maintaining a certain level of quality of service for ongoing connections. Admission control is closely coupled with other resource allocation schemes, such as dynamic channel allocation, power control, and mobility prediction — to name a few. Furthermore, admission control becomes more challenging in the content of supporting multimedia services with different and multi-faceted QoS requirements in a wireless environment. One could take two approaches to the admission control problem: (1) to admit a user based on whether the network can satisfy the QoS requirements of all the users at the time of admission; (2) to admit a user based on whether the network can satisfy the QoS requirements of all users taking into consideration mobility and variation in the channel performance. A solution that only considers the former has the appeal of simplicity, but the admissible decision would inherently have to be conservative (this approach is used, for example, with channel reservation strategies for handoffs in circuit-switched cellular systems [14]). The latter approach would have the advantage of more intelligent decision-making at the cost of increased computation [15].

Multiple access

Multiple access techniques allow a communication medium to be shared among different users. The three basic multiple access techniques are frequency-division multiple access (FDMA), time-division multiple access (TDMA), and code division multiple access (CDMA). First generation analog cellular systems use FDMA. Both TDMA and CDMA techniques are implemented in second generation digital cellular systems and are competing for third generation standards [16]. The third generation CDMA standard in North American is cdma2000 while the WCDMA standard is specified in Europe and Japan. Enhanced Data Rates for GSM Evolution (EDGE) standards are currently specified in order to provide a third-generation evolution option for TDMA systems. The two components of EDGE, enhanced circuit switched data (ECSD) and EGPRS, define enhancements for circuit-mode and packet-mode data, respectively.

OFDM

Current cellular networks provide satisfactory voice services at a reasonable cost. However, this success story does not easily extend to data services, such as wireless web surfing. Whether wireless channels can provide high-data-rate service and whether such a service is affordable are two major concerns. To provide high-data-rate service, wide-band transmission is necessary. In a wide-band single-carrier system, we face the problems of frequency-selective-fading and inter-symbol-interference. Furthermore, to make high-rate-data service affordable, a higher spectrum efficiency (compared with current systems) has to be achieved. OFDM (orthogonal frequency division multiplexing) is a promising transmission (modulation) technique [17, 18, 19, 20] to combat ISI over multipath fading channels and provide efficient frequency utilization. This technique appears to show promise for high-speed wireless/wireline data communications. A properly coded and interleaved OFDM system is reported to exceed the performance of many other existing systems. For example, Flarion claims that their flash-OFDM airlink enables three times higher spectral efficiency than the CDMA 3G airlink, and the cost per Megabyte of data is about 10.5 cents [21, 22].

1.2 Adaptation Techniques

Today's cellular systems are designed to provide good coverage for voice services, providing the required minimal data rate everywhere, that is, to achieve a minimum required signal to interference plus noise ratio (SINR) over 90–95 percent of the coverage area. This ensures that the data rate required to achieve "good" voice quality can be provided "everywhere." As a result, SINRs that are much larger than the target minimum are achieved over a large portion of the cellular coverage area.

For packet data service, the larger SINRs can be used to provide higher data rates by reducing coding or spreading and/or increasing the constellation density. Research shows that cellular spectral efficiency (in terms of b/s/Hz/sector) can be increased by a factor of two or more if users with better links are served at higher data rates [23]. Procedures that exploit this are already in place for the major cellular standards in the world.

Rate adaptation in CDMA systems is achieved through a combination of variable spreading, coding, and code aggregation. Wideband CDMA (WCDMA) and cdma2000 systems achieve higher rates through a combination of variable spreading and coding. The WCDMA standard, being specified in Europe and Japan, supports data rates up to 2.048 Mb/s in 5 MHz bands (using 6 code channels simultaneously). The third generation North American CDMA standard (cdma2000) supports a data rate up to 614.4kb/s in 5 MHz bands with the lowest spreading factor of 2.

In TDMA systems, adaptive coding, adaptive modulation, and incremental redundancy are used to achieve variable data rates. GPRS-136 and EGPRS specify the use of incremental redundancy transmission. GPRS-136 employs adaptive modulation and incremental redundancy to achieve higher throughput. EGPRS, the enhanced packet-mode data standard for a 3G evolution option for TDMA system, exploits adaptive coding, adaptive modulation, and incremental redundancy to achieve a peak rate of 473.6kb/s (when 8 time-slots are aggregated).

For incremental redundancy transmission, an initial transmission of data is uncoded or lightly coded. On decoding failure at the receiver, incremental amounts of redundancy are transmitted until the receiver is able to successfully decode the data frame. For example, if a rate-1/2 non-systematic convolutional code is used, the output of one generator is mapped to the D data blocks, and the output of the other generator is mapped to the D parity blocks. The D data blocks are transmitted first, resulting in a code rate of unity for the first transmission. In response to each negative acknowledgment, one parity block is transmitted to achieve the code rates: $D/(D+1), \cdots, 1/2$. Thus, the procedure effectively matches the coding rate to the channel SINR without requiring SINR estimation and feedback. In addition, the transmission of redundant information dispersed in time provides a diversity advantage during decoding. Incremental redundancy is also being considered for WCDMA standards.

With rate adaption techniques, these packet data standards mentioned before achieve cellular spectral efficiency in the range 0.1–0.3 b/s/Hz/sector, while in second generation cellular systems, the voice services achieve cellular spectral efficiencies in the range of 0.03–0.05 b/s/Hz/sector.

Research shows that fast rate adaptation is required to achieve high capacity on fast fading channels [24]. In IS-856 [25, 26], the pilot bursts provide the mobile users with the means to estimate accurately and rapidly the channel conditions. Among other parameters, each mobile user estimates the received E_c/N_t of all resolvable *multipath* components and predicts the effective received SINR. This channel state information is then fed back to the base station via the reverse link data rate request channel (DRC) and updated as quickly as every *1.67 ms*.

Channel condition feedback is important to rate adaptation [23]. In CDMA systems, pilot strength measurements are used to estimate the SINR at the receiver. In IS-95B and cdma2000, pilot strength measurements are provided to the base station through the pilot strength measurement message (PSMM) or included in the supplemental channel request message (SCRM). The measurement report message in WCDMA can additionally include block error rate, bit error rate (BER), received power, path loss, and downlink SINR measurements. In TDMA systems, channel quality is estimated at the receiver and the information is provided to the transmitter through appropriately defined messages. In both GPRS and EGPRS systems, measurement reports are included in supervisory ARQ status messages. Metrics that have been proposed for estimating channel quality are: frame error rate, mean and standard deviation of symbol error rate (SER) or BER, average SINR. The GPRS measurement reports consist of an estimated BER and a variance of the BER estimated over a short moving window. In GPRS-136, the receiver provides channel quality feedback as an indication of the modulation schemes that are permitted based on the estimated SINR.

In summary, adaptation techniques have been developed so that a user can achieve high performance when it is in a good condition. Based on this technique, transmission scheduling further improves spectrum efficiency by letting users transmit in relatively good conditions, as we will discuss in the following section.

1.3 Transmission Scheduling in Wireless Networks

In wireline networks, resource allocation schemes and scheduling policies play important roles in providing service performance guarantees, such as throughput, delay, delay-jitter, fairness, and loss rate [27]. Scheduling disciplines and associated performance problems have been widely studied in packet-switched networks [28, 29]. There are two types of scheduling disciplines: work-conserving and non-work-conserving. A server using a work-conserving discipline is never idle when there is a packet to be sent. With a non-work-conserving discipline, each packet is assigned, either explicitly or implicitly, an eligibility time. Even when the server is idle, if no packets are eligible, none will be transmitted. Examples of work-conserving scheduling disciplines are: Delay Earliest-Due-Date (Delay-EDD), Virtual Clock, Fair Queuing (FQ), Weighted Fair Queuing (WFQ), and Worst-case Fair Weighted Fair Queueing (WF^{2}Q). However, with work-conserving disciplines, the traffic pattern is distorted inside the network due to fluctuations in the network load. For services that require guaranteed performance, the more important performance index is the end-to-end delay bound rather than the average delay. This is one of the motivations for non-work-conserving scheduling policies. Several non-work-conserving disciplines have been proposed [27]: Jitter Earliest-Due-Date (Jitter-EDD), Stop-and-Go, Hierarchical Round Robin (HRR), and Rate-Controlled Static Priority (RCSP). In addition to the challenge of providing service performance guarantees, scheduling disciplines must be *simple* and *scalable* to be implemented in real networks due to the size of wireline networks.

It is important to note that resource allocation and scheduling schemes from the wireline domain do not carry over to wireless systems because wireless channels have unique characteristics not found in wireline channels. Some of these characteristics are:

- Channel conditions are constantly varying.
- Network performance depends on channel conditions and signal processing techniques.
- If the same resource is given to different users, the resultant network performance (e.g., throughput) could be different from user to user.

We next explain these characteristics in more detail and why they are important to the design of wireless scheduling policies. In wireless networks, the channel conditions of mobile users are time-varying. Radio propagation can be roughly characterized by three nearly independent phenomena: path-loss variation, slow log-normal shadowing, and fast multipath-fading. Path losses vary with the movement of mobile stations. Slow log-normal shadowing and fast multipath-fading are time-varying with different time-scales. Furthermore, a user receives interference from other transmissions, which is time-varying; and background noise is also constantly varying. Hence, mobile users perceive time-varying channel conditions. SINR (signal to interference plus noise ratio) is a commonly used measure of channel conditions. Fig. 1.1 shows the time-varying SINR of a mobile user. Other measures include BER (Bit Error Rate) and FER (Frame Error Rate).

Because channel conditions are time-varying, users experience time-varying service quality and/or quantity. For voice users, better channel conditions may result in better voice quality. For packet data service, better channel conditions (or higher SINR) can be used to provide higher data rates using adaption techniques explained in Section 1.2. Research shows that cellular spectral efficiency (in terms of b/s/Hz/sector)

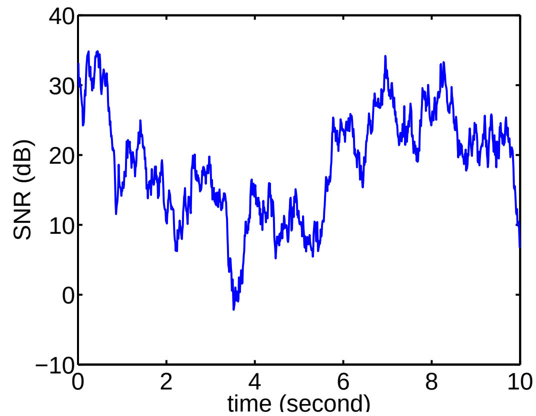

Fig. 1.1. User's time-varying SINR

can be increased by a factor of two or more if users with better links are served at higher data rates [23]. Procedures to exploit this are already in place for all the major cellular standards: adaptive modulation and coding schemes are implemented in the TDMA standards, and variable spreading and coding are implemented in the CDMA standards. In general, a user is served with better quality and/or a higher data rate when the channel condition is better. Hence, good scheduling schemes should be able to exploit the variability of channel conditions to achieve higher utilization of wireless resources.

The performance (e.g., throughput[1]) of a user depends on the channel condition it experiences, hence, we expect different performance when the same resource (e.g., radio frequency) is assigned to different users. For example, consider a cell with two users. Suppose that user 1 has a good channel, e.g., it is close to the base station. User 2 is at the edge of the cell, where the path-loss is significant and the user experiences large interference from adjacent cells. If the same amount of resource (power, time-slots, etc.) is assigned, it is likely that the throughput of user 1 will be much larger than that of user 2.

[1]Throughput is the number of useful information bits successfully transmitted between the base station and mobile users during a unit time.

Different assignments of the wireless resource will affect the system performance, hence, resource allocation and scheduling policies are critical in wireless networks. In the dissertation, we study *opportunistic* scheduling; by opportunistic, we mean the ability to exploit the variation of channel conditions. Earlier, we have described the unique features of wireless networks. Then an important question is: under such conditions, what should be the basic features of a scheduling policy? Consider a few users that share the same resource. The users have constantly varying channel conditions, which imply constantly varying performance. The scheduling policy decides which user should transmit during a given time interval. Intuitively, we want to assign resource to users experiencing "good" channel conditions so that the resource can be used efficiently. At the same time, we also want to provide some form of fairness or QoS guarantees to all users. For example, allowing only users close to the base station to transmit with high transmission power may result in very high system throughput, but may starve other users. This basic dilemma motivates our work: to improve wireless resource efficiency by exploiting time-varying channel conditions while at the same time control the level of fairness/QoS among users.

Fairness criteria may have different implications in wireline and wireless networks. In wireline networks, when a certain amount of resource is assigned to a user, it is equivalent to granting the user a certain amount of throughput/performance value. However, the situation is different in wireless networks, where the amount of resource and the performance values are not directly related (though correlated). Hence, we study two kinds of fairness: temporal vs. utilitarian. Temporal fairness means that each user gets a fair share of network resource, and utilitarian fairness means that each user gets a certain share of the overall system capacity. Further, we consider both long-term fairness and short-term fairness in the dissertation. The basic idea of opportunistic scheduling is to let users transmit in "good" channel conditions, and thus a natural question is how long a user is willing to wait (for good conditions). Hence, there exists tradeoff between scheduling performance gain and short-term performance. In addition to fairness, we consider a long-term QoS metric: each user has a specific data-rate requirement for the system. Because the capacity of a wireless system is not fixed, it is not always an easy task to determine the feasibility of the requirements of all users. We will study these issues in more detail in the dissertation.

Interference management is also a crucial component of efficient spectrum utilization in wireless systems because interference limits the system capacity ultimately. Power allocation is a traditional interference management mechanism. It has been well studied and widely used in wireless systems to maintain desired link quality, minimize power consumption, and alleviate interference to others [8]. Further, because users experience time-varying and location-dependent channel conditions in wireless environments, we can schedule users "opportunistically" so that a user can exploit more of its good channel conditions and avoid (as far as possible) bad times, at least for applications (e.g., data service) that are not time-critical. Hence, joint scheduling and power-allocation scheme should be able to further improve the spectrum efficiency and decrease power consumption.

Opportunistic scheduling exploits the variation of channel conditions, and thus provides an additional degree of freedom in the time domain. Moreover, it can be coupled with other resource management mechanisms to further increase network performance. A good example of it is joint scheduling and power allocation as explained earlier. In the literature, opportunistic scheduling is also referred as multiuser diversity [30]. Occasionally, these two terms may have slightly different meanings. An example is the case where there is only one user in the system and the objective is to minimize transmission power while maintaining a certain data rate.

1.4 Scope of the Report

We now briefly describe the organization of the dissertation. In Chapter 2, we describe the system model used in the dissertation and review some related work. Chapter 3 provides a case study of opportunistic scheduling with temporal fairness. Here, we present an opportunistic transmission scheduling policy that exploits time-varying channel conditions and maximizes the system performance stochastically under the temporal fairness constraint. We prove the optimality of the scheduling scheme with the assumption that the system is stationary. We also prove that every user improves its performance over a non-opportunistic scheduling policy when users have independent performance values. Through simulation results, we show that the scheme also works well for non-stationary scenarios and results in significant performance improvements over a scheduling scheme that does not take into account channel conditions. We also show that our scheduling scheme is robust to estimation errors. Further, we describe a scheme to improve "short-term" performance. In Chapter 4, we present a unified *framework* for opportunistically scheduling user transmissions to exploit the time-varying channel conditions in wireless communication systems. The objective is to maximize the system performance while satisfying various QoS requirements. Our framework enables us to investigate three categories of scheduling problems involving two fairness requirements (*temporal fairness* and *utilitarian fairness*) and a minimum-performance requirement. We provide optimal scheduling solutions, discuss the advantages and disadvantages of the various scheduling formulations, and study the asymptotic behavior of our opportunistic scheduling schemes. We also show that our results can be generalized to include non-stationary policies in a non-stationary environment. In Chapter 5, we present joint scheduling and power-allocation schemes to alleviate intercell interference. First, we study the problem with the objective to minimize the average transmission power, and thus interference to other cells, while maintaining the required data-rate for each user within the cell. Then we study the problem to maximize the net utility, defined as the difference between the value of throughput and the cost of power consumption, with the same data-rate requirements. We establish the optimality of our joint scheduling and power-allocation schemes for both problems. Finally, we conclude our work and outline proposals for future work in Chapter 6.

2. SYSTEM MODEL AND RELATED WORK

2.1 Land-Mobile Radio Propagation

The three basic propagation mechanisms which impact propagation in a mobile communication system are reflection, diffraction, and scattering [31]. Reflection occurs when a propagating electro-magnetic wave impinges upon an object that has very large dimensions compared to the wavelength of the propagating wave. Reflections occur from the surface of the earth and from buildings and walls. Diffraction occurs when the radio path between the transmitter and receiver is obstructed by a surface that has sharp irregularities (edges). The secondary waves resulting from the obstacle are present throughout the space and even behind the obstacle, giving rise to a bending of waves around the obstacle. Scattering occurs when a medium through which the wave travels consists of objects with dimensions that are small compared to the wavelength, and where the number of obstacles per unit volume is large. Scattered waves are produces by rough surfaces, small objects, or by other irregularities in the channel. In practice, foliage, street signs, and lamp posts induce scattering in a mobile communication system.

We next describe the propagation models that reflect the impact of these three basic propagation mechanisms. Land-mobile radio signal can be characterized in terms of three propagation regimes:

- power-law propagation;
- log-normal shadowing;
- fast fading.

Power-law propagation is on the largest scale. Both theoretical and measurement-based propagation models indicate that the average received signal power decreases logarithmically with distance, whether in outdoor or indoor environments. Such models have been used extensively in the literature. The average large-scale path loss for an arbitrary transmitter-receiver separation is expressed as a function of distance by using a path loss exponent. For example, in Lee's model [32], the path loss l_p (dB) is

$$l_p = K + 10\alpha \log_{10}(d) - \alpha_0,$$

where d is the distance between the transmitter and receiver, α is the path loss factor, α_0 is a correction factor used to account for different base station and mobile station (MS) antenna heights, transmit powers, and antenna gains, and K is a constant, which has different values in different environments.

Furthermore, received power variations of many decibels can occur due to shadowing even for small shifts in range. An explanation for lognormal shadowing is as follows. Consider the received signal to be the result of the transmitted signal passing through or reflecting off some random number of objects such as buildings, hills, and trees. The individual processes each attenuate the signal to some degree and the final received value is thus the product of many transmission efficiency factors. Therefore, the logarithm of the received signal equals the sum of a large number of factors, each also expressed in decibels. As the number of factors becomes large, the central limit theorem shows that the distribution of the sum can be modeled as a Gaussian distribution under fairly general assumptions [33]. The shadowing term $s(k)$ (dB) is modeled as a zero-mean stationary Gaussian process with autocorrelation function given by

$$E(s(k)s(k+m)) = \sigma_o^2 \xi_d^{vT/D},$$

where ξ_d is the correlation between two points separated by a spatial distance D (meters), and v is velocity of the mobile user. In our simulation, we use a value of $\sigma_o = 4.3$ dB, corresponding to a correlation of 0.3 at a distance of 10 meters, as reported by Gudmundson [34].

Finally, as a mobile moves over very small distances, the instantaneous received signal strength may fluctuate rapidly giving rise to small-scale fading. The reason for this is that the received signal is sum of many contributions coming from different directions because local scatterers commonly surround the receiver antenna. Since the phases are random, the sum of the contributions varies widely; for example, obeys a Rayleigh fading distribution. During small-scale fading, the received signal power may vary by as much as three or four orders of magnitude when the receiver is moved by only a fraction of a wavelength [31].

In summary, radio propagation can be roughly characterized by three nearly independent phenomena: path-loss variation, slow log-normal shadowing, and fast multipath-fading. Path losses vary with the movement of mobile stations. Slow log-normal shadowing and fast multipath-fading are time-varying with different time-scales. Also the interference a user received due to other transmissions is time-varying. Furthermore, background noise is also constantly varying. All this contribute to the channel having time-varying characteristics and motivates the need for "opportunistic" scheduling schemes.

2.2 Literature Review

In wireline networks, resource allocation schemes and scheduling policies play important roles in providing service performance guarantees, such as throughput, delay, delay-jitter, fairness, and loss rate. Scheduling disciplines and associated performance problems have been widely studied in packet-switching networks [28, 29]. For a good survey of such algorithms, see [27]. However, as mentioned in the introduction, scheduling schemes from the wireline domain do not carry over to wireless systems because wireless channels have unique characteristics not found in wireline channels.

Transmission scheduling for wireless networks has recently attracted a lot of attention. First, scheduling policies of wireline networks are extended to wireless networks, where the bursts of errors in wireless channels is taken into account. To elaborate, a wireless channel is modeled by a two-state Markov chain [35]: 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 [36, 37, 38, 39]. 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 [40]. 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 has been developed at Qualcomm to provide a versatile wireless Internet solution [26]. This system is also known as High Data Rate (HDR) [25]. The first fundamental design choice of HDR is to separate the services by including two interoperatble modes: that is, 1x mode for voice and low-rate data and 1xEV mode for high-rate data services. In 1xEV mode, a single user is served at any instant (e.g, time-multiplexed CDMA); therefore avoiding power sharing and allocating the entire access point (e.g., base station) power to the user being served. The Access Point always transmit at full power achieving very high peak rates for users that are in a good coverage area. The Access Terminal, on a slot-by-slot basis (1.67 ms), measures the pilot strength, and continuously requests an appropriate data rate based on the channel conditions.

The scheduler in IS-856 uses a different notion of fairness known as proportional fairness [41, 42]. 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 $x\%$. To be specific, assume that there are N users and $R_i(t)$ is the estimate of the average rate for user i for slot t, $i = 1, \cdots, N$. Also, suppose that at slot t, the current DRC (i.e., requested date rate) from user i is $DRC_i(t)$, again $i = 1, \cdots, N$. The algorithm works as follows:

- Scheduling: The user with the highest ratio of $DRC_i(t)/R_i(t)$ of all N users will receive transmission at each decision time. Ties are broken randomly.

- Update Average Rate: For each user i,

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

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 [43].

The author of [44, 45] analyzes the PF scheduler under simplified conditions. Let

$$(C/I)_i(t) = a_i b_i(t),$$

where a_i is the distance dependent component of C/I for user i and $b_i(t)$ is the random component of C/I for user i, including Rayleigh or Ricean fading. The author of [44, 45] assumes that the rate is a linear function of $C/I(t)$ and $b_i(t)$s are independent among the users and across time. It is then shown that when $b_i(t)$ are i.i.d. (independent and identically distributed) among users, the PF scheduler gives equal power and time to users and the throughput of individual users, T_i, is inversely proportional to a_i; i.e., T_i/a_i equals a constant value. Further, when users have different distributions of $b_i(t)$ (such as Rayleigh vs. Ricean), the user with the higher variability of $b_i(t)$ gains higher throughput while using a (slightly) smaller amount of time.

In [46, 47, 48], 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 channel capacity, and ρ_i is some constant. The proposed scheduler achieves throughput optimality, defined in [46] 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, the authors of [47] state the following result (assuming there is a finite set of channel states): to maximize the system throughput with minimum-throughput requirements, there exists some constant c_i such that one should choose a user with the maximum value of $c_i R_i$. In these papers, however, there is no discussion on how to obtain the values of c_i, how to break ties, or how feasibility plays a role. Further, in [49, 50], the authors study an exponential rule:

$$\operatorname{argmax} \rho_i R_i \exp\left(\frac{a_i W_i}{\beta + \bar{W}^{\eta}}\right),$$

where W_i is the queue length (or waiting time), and $\bar{W}$ is the average queue length (or waiting time) over users. Using the exponential rule, when all queues are filled to similar capacity, the channel condition plays a significant part. On the other hand, if one queue is much longer than others, then the queue length becomes dominant and the longer queue gets a higher chance to transmit. Hence, this algorithm balances the tradeoff between queue length and throughput. The exponential rule is also throughput optimal.

The authors of [51] investigate a scheduling algorithm to maximize the minimum (weighted) throughput of users. To elaborate, the objective function is

$$\text{maximize} \quad \min_i \lim_{N\to\infty} E\left(\sum_{t=1}^{N} \mathbf{1}_{\{\}} i R_i(t)\right),$$

where $R_i(t)$ is the rate of user i at time t, $\mathbf{1}_{\{\}} i = 1$ if time-slot t is assigned to user i, and $\mathbf{1}_{\{\}} i = 0$ otherwise. The 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.

In [52], 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 [52], the tradeoffs between the required power and various notions of delay are analyzed.

The authors of [12] study joint power control and intracell scheduling in DS-CDMA systems. Assuming that the data rate is a linear function of its SINR, it shows that scheduling users one-by-one within a cell results in better performance than simultaneous transmission within a cell, especially when the data rate requirement is high. Hence, via one-by-one intracell transmission, the required transmission power of a cell is minimized. In addition to the intracell one-by-one scheduling, they suggest a distributed power control scheme for intercell interference management. Basically, each base station calculates the minimum power needed to support the required data rates of users in the cell. The transmission power remains constant within each scheduling interval. The base stations then update their required powers based on interference from other cells without intercell communication, which is similar to standard distributed power control algorithms. It is shown that the proposed distributed power control algorithm converges, the rate of convergence is geometric, and that the resulting power is the minimum required to support the required data rate. However, this work in its current form does not exploit time-varying property of channels; i.e., no opportunistic scheduling.

In [53], the authors study scheduling problems for real-time traffic with fixed deadlines. Scheduling in a time-slotted system is considered, the capacity of the channel is time-varying, and the BS can estimate the channel of the current time-slot. The mobiles achieve different QoS based on the unit prices that they are willing to pay. The objective of the base station is to maximize the revenue of the base station. The scheduling is preemptive and the base station obtains a partial revenue if a request is served partially. The unit price of a request is a non-increasing function of the time. The offline optimal scheduling scheme is shown to be NP-complete if only one user can be assigned to a time-slot. The authors then propose a greedy algorithm that chooses the request with the largest revenue in the current time-slot to serve. The authors show that the greedy algorithm is $1/2$ competitive against the offline optimal algorithm. Further, they show that no deterministic online algorithm can achieve a competitive ratio higher than $1/2$. (This does not mean that the greedy algorithm will always do better than other deterministic online algorithms.) Then the authors extend the work to various scenarios such as multi-carrier case, the case where a single slot can be shared by several users, and the case where the price is a non-increasing function of the total data that has been served to this request (which may be applicable to layered multimedia data where base layer is more important than enhanced layer).

Downlink scheduling in CDMA systems for data transmission is also studied in [54]. The work considers a performance metric called "stretch", which is defined as the delay experienced by a packet normalized by its minimum achievable delay. The stretch can be considered as 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 [55], the authors study transmission schemes for time-varying wireless channels with partial state information. A finite-state Markov chain is used to model the channel, and channel information is only available at the end of the time-slot if the transmission occurs during the time-slot. It is assumed the channel transmission matrix is known. The objective is to minimize a discounted infinite-horizon cost function, which can be used to indicate the balance between power cost and throughput. An example of the cost function is:

$$C(g,s) = \begin{cases} c_0 & s = 0 \\ c_1 s + c_2 P_e(gs) & s > 0, \end{cases}$$

where g is the state, s is the transmission power, and P_e is the error probability. The resulting optimal solution is a threshold back-off scheme: suppose a packet transmission occurs during the last time-slot and the channel state is known. If the current minimum cost is greater than c_0 (no transmission cost, penalizes a scheme for placing too much emphasis on energy efficiency), then the system keeps silence for a certain number of time-slots, and then resumes transmission. The optimal transmission power is the power that minimizes the current cost function. The paper studies the effect of channel memory with partial state information, while the result may depend on the accuracy of the POMDP (Partially Observable Markov Decision Process) channel models and the estimation of transmission matrix.

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 [30], the authors use multiple transmission antennas to "induce" channel fluctuations, and thus exploit multi-user diversity. Consider a static channel (static in the time-scale of interest) and N multiple transmission antennas. Let $h_{ni}(t)$ be the channel gain from antenna n to user i at time t. At time t, $x(t)$ is multipled by $\sqrt{a_n(t)}e^{j\theta_n(t)}$ and transmitted at antenna n, $i = 1, \cdots, N$, where $\sum_{n=1}^{N} a_n(t) = 1$ to preserve the total transmission power. Here, $a_n(t)$ and $\theta_n(t)$ are random variables used 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. Note that the optimal beam-forming configuration is:

$$\begin{aligned} a_n &= \frac{|h_{ni}|^2}{\sum_{n=1}^{N} |h_{ni}|^2}, \quad n = 1, \cdots, N \\ \theta_n &= -\arg h_{ni}, \quad n = 1, \cdots, N, \end{aligned}$$

which requires individual channel information (amplitude and phase) from each antenna.

2.3 System Model

The current PCNs (Personal Communication Networks) use a cellular architecture. The geographical coverage area is partitioned into cells, each served by a base station. Mobile users are connected to the network via the base stations. Within a cell, users share resource in terms of time, frequency, power, and/or code.

In the dissertation, we consider a time-slotted system—time is the resource to be shared among all users. At any given time, only one user can occupy a given channel (frequency band) in a cell; multiple users could transmit at different frequencies within a cell. We focus on the scheduling problem for one channel. Note that a channel in this context could be very large. For example, it is possible for 10 users to share a 1MHz frequency band for high rate data service, while in the IS-136 standard, a voice channel uses 10KHz bandwidth.

In a time-slotted system, the access technique can either be TDMA (Time Division Multiple Access) or CDMA (Code Division Multiple Access). It is obvious that our system model fits TDMA systems. The time-slot in the system model is a natural match for the time-slot in a TDMA system. Note that the time-slot in the system model could consist of more than one contiguous time-slot in the TDMA system.

The system model also fits the need for data traffic in CDMA systems. In traditional CDMA systems, (e.g., IS-95), a multitude of low-data-rate channels are multiplexed together (with transmissions made orthogonal in the code domain) and share the available base station transmitted power with some form of power control. This is an optimal choice for many low-rate channels sharing a common bandwidth. The situation becomes less optimal when a small number of high-rate users share the channel. The inefficiencies increase further when the same bandwidth is shared between low-rate and high-rate data users, since their requirements are vastly different [25]. Research shows that, to achieve high-data capacity, data users should transmit in a time-multiplexed mode instead of transmitting simultaneously [56, 57].

In the IS-856 standard, the downlink is designed differently from that in IS-95. The downlink transmissions in IS-856 are time multiplexed and transmitted at the full power available to the mobile users [26]. This means that a single user is served at any instant with full transmission power. Hence, the system model in the dissertation also fits time-slotted CDMA systems, which is a potential solution for providing high-data-rate service in CDMA systems.

The scheduling scheme can be applied to both downlink and uplink. In general, downlink transmission is more important for data traffic due to the highly asymmetric nature of the data service. Further, the uplink may experience synchronization difficulties due to different distances between users and the base station when the duration of a time-slot is short.

As explained previously, channel conditions in wireless networks are time varying and thus users experience time-varying performance. Throughout the dissertation, we use SINR (signal to interference and noise ratio) as a measure of channel conditions [58, 59]. Note that SINR is related to E_b/N_0 (the ratio of bit energy to noise power spectrum density) as

$$SINR = \frac{E_b R}{N_0 W},$$

where W is the channel bandwidth and R is the bit rate. We use a stochastic model to capture the time-varying channel condition of each user. To elaborate, let $\{\alpha_i^k\}$ be a stochastic process associated with user i, where α_i^k represents the received SINR for user i at time k given that the transmission power is 1. For analytical simplicity, we assume that the stochastic process $\{\alpha_i^k, k = 1, 2, \cdots\}$ is stationary and ergodic. Hence, we drop the time index k.

In Chapters 3 and 4, we use an abstract concept to measure *time-varying* and *channel-condition-dependent* performance of each user. Specifically, let U_i be a random variable associated with user i, measuring the level of performance that would be experienced by user i if it is scheduled to transmit at a generic time-slot. The value of U_i measures the "worth" of a time-slot to the user i, and is in general a function of its channel condition, indicated by SINR. The value of U_i could also depend on other factors, such as user i's coding/modulation scheme and power consumption. Usually, the better the channel condition of user i, the larger the value of U_i. Denote $\vec{U} = (U_1, \cdots, U_N)$, where N is the number of users.

Recently, the notion of "utility" has been used to study wireless resource management problems [60, 61]. Utility is generally defined as a measure of satisfaction that a user derives from accessing the wireless resource (in terms of bandwidth, power, etc). Along these lines, we can think of the value of U_i^k as the "utility value" of the channel to user i at time k. However, unlike the typical usage of utility, note that in the dissertation, we model a user's channel-utility value as a stochastic process, capturing the important feature of wireless systems that channel conditions are time-varying.

We next present examples of possible performance measures. The most straightforward performance measure is the throughput (in terms of bits/sec) or the "monetary value" of the throughput (in terms of dollars/sec), where the throughput is the number of information bits per time-slot successfully transmitted between the base station and the mobile user. Usually, a user's throughput is a nondecreasing function of SINR (signal to interference and noise ratio). Depending on the class of a user, the throughput could be a step function, an S-shape function, or a linear function of the SINR, as shown in Figure 2.1. Hence, different classes of users may have different throughput values even with the same channel condition. In our system model, we do not make any assumption on the physical-layer implementation of the system. Note also that the throughput of a user could be limited by the user's interference to other cells. For example, consider a user at the edge of a cell, where the user's transmission causes significant interference to a neighboring cell. When the neighboring cell is heavily loaded, the user's maximum transmission power may have to be limited to avoid undue interference to the neighboring cell.

Besides throughput, other issues could also be important to users and different users could have different utility functions. For example, a user on a vehicle where there is no scarcity of power may be concerned only about the throughput. On the other hand, power consumption is very important to a handset user, and hence the performance of such a user could have the form:

$$\text{value of throughput} - \text{cost of power consumption.}$$

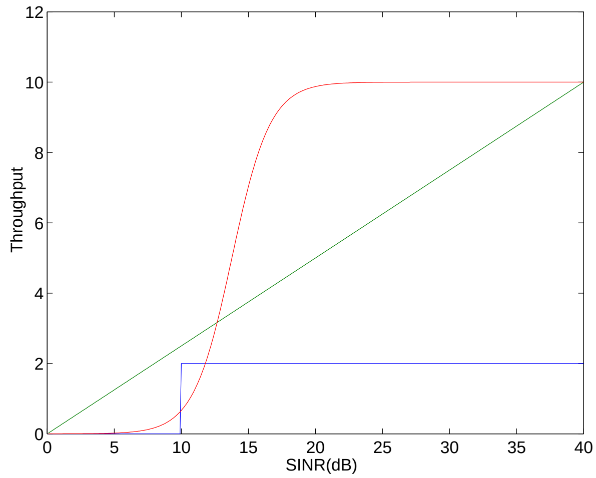

Fig. 2.1. Users' throughput as a function of SINR.

In summary, the performance value U_i is an abstraction used to capture the time-varying and channel-condition-dependent "worth" of a time-slot to a user. The use of such a general model frees us from physical-layer implementation details and allows us to focus on the problem of designing scheduling policies. We assume throughout that performance values for different users are *comparable and additive*. For example, the unit of the performance values could be "dollars per time unit" for all users. In other words, if $U_1^k > U_2^k$, then assigning time-slot k to user 1 rather than user 2 will lead to higher system performance. Also, if U_1^1 and U_1^2 are the performance values for user 1 at time-slots 1 and 2, then the total performance over the two time-slots is $U_1^1 + U_1^2$.

In this chapter, we describe the propagation model and system model. In the following chapters, we will use the system model to study opportunistic scheduling problems.

3. TEMPORAL FAIRNESS SCHEDULING: A CASE STUDY

In this chapter, we study an opportunistic scheduling problem with temporal fairness constraints based on the system model presented in Section 2.3. First, we present an opportunistic scheduling policy that maximizes the system performance stochastically under the temporal fairness constraint. Then we describe a stochastic-approximation-based algorithm that can be used to efficiently estimate the key parameters of the scheduling scheme on-line. Through simulation results, we show that the scheme also works well for non-stationary scenarios, is robust to estimation errors, and results in significant performance improvements over a scheduling scheme that does not take into account channel conditions. Last, we discuss a scheme to improve "short-term" performance.

3.1 Problem Formulation

Because time is the resource shared among users, a natural fairness criterion is to give each user at least a certain share of the system resource, i.e., time. The temporal fairness requirement is defined as follows: suppose that there are N users in a cell, each user i is assigned a fixed fraction of resource (i.e., time-slots), denoted as r_i, where $0 \leq r_i \leq 1$ and $\sum_{i=0}^{N} = 1$. That means, on average, r_i portion of time-slots should be used by user i. Note that r_i is predetermined by time-slot assignment algorithm, whose value typically depends on the user's class, the price the user is willing to pay for the wireless service, or the user's current channel conditions.

Recall that $\vec{U} = (U_1, \cdots, U_N)$, where U_i is a random variable representing the performance value of user i at a generic time-slot as explained in Section 2.3. The scheduling problem is stated as follows: given $\vec{U}$, determine which user should be scheduled (in the given time-slot). We define a *policy* Q to be a mapping from the performance-vector space to the index set $\{1, 2, \cdots, N\}$. Given $\vec{U}$, the policy Q determines the user to be scheduled: if $Q(\vec{U}) = i$, then user i should use the time-slot, and the system receives a performance "reward" of $U_{Q(\vec{U})}$ (i.e, U_i). Hence, $E(U_{Q(\vec{U})})$ is the average system performance value associated with policy Q. Note that the policy Q is potentially "opportunistic" in the sense that it can use information on the performance vector $\vec{U}$ to decide which user to schedule.

We are interested only in policies that result in satisfaction of the temporal fairness constraints. Specifically, we say that a policy Q is *feasible* if $P\{Q(\vec{U}) = i\} = r_i$ for all $i = 1, \cdots, N$. Feasible policies are those that obey the given fairness constraints. We use Θ to denote the set of all feasible policies.

Our goal is to find a feasible policy Q that maximizes the average system performance while satisfying the fairness constraints. The problem can be stated formally as follows:

$$\underset{Q \in \Theta}{\text{maximize}}\, E\left(U_{Q(\vec{U})}\right). \tag{3.1}$$

Note that we can write

$$\begin{aligned} E\left(U_{Q(\vec{U})}\right) &= E\left(\sum_{i=1}^{N} U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right) \\ &= \sum_{i=1}^{N} E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right), \end{aligned}$$

where

$$\mathbf{1}_A = \begin{cases} 1 & \text{if } A \text{ occurs,} \\ 0 & \text{otherwise} \end{cases}$$

is the indicator function of the event A. In other words, the overall objective function is the sum of all users' average performance values (where we reap a reward of U_i only if user i is scheduled).

Recall that we assumed the sequence $\{\vec{U}^k\}$ to be stationary. This assumption does not preclude correlations across users or across time. In practice, a user's channel condition is usually time-correlated, for example, due to shadowing. Hence, a user's performance is usually also time-correlated. Furthermore, the performance of different users may also be correlated. For example, when the intercell interference is high, most users' performance values simultaneously decrease. However, if users have enough separated locations, it is reasonable to assume that their performance values are only weakly dependent.

We first present our opportunistic scheduling policy given the temporal fairness constraints. (We provide the proof of its optimality in Appendix A.1.) We then explain how to estimate the parameters used in the policy. Further, we describe a procedure to implement our scheduling policy by tuning the parameter values based on measurements, and provide simulation results. Finally, we study a heuristic scheduling algorithm to improve short-term performance.

3.2 Opportunistic Scheduling Policy

3.2.1 2-user Case

For the purpose of illustration, we start with the 2-user case. Suppose that user 1 and user 2 have temporal requirements r_1 and r_2, respectively, and $r_1 + r_2 = 1$. We wish to find an opportunistic policy that solves (3.1).

Define $y(v) = P\{U_1 + v \geq U_2\}$, where $v \in \mathbb{R}$. Because $y(v)$ is the distribution function of the random variable $(U_2 - U_1)$, $y(v)$ is a right-continuous monotonically increasing function of v with $y(\infty) = 1$ and $y(-\infty) = 0$. Hence, there exists a v^* (which may not be unique) such that for any $\epsilon > 0$,

$$y(v^* - \epsilon) \leq r_1 \leq y(v^*),$$

where r_1 is the temporal requirement of user 1.

We consider the scheduling policy under two conditions.

1. $y(v^*) = r_1$: The opportunistic scheduling policy in this case is given by

$$Q^*(\vec{U}) = \begin{cases} 1 & \text{if } U_1 + v^* \geq U_2, \\ 2 & \text{otherwise.} \end{cases}$$

2. $y(v^*) > r_1$: Let $y^-(v) = P\{U_1 + v > U_2\}$, which is a left-continuous monotonically increasing function of v. So

$$y^-(v^*) \leq r_1 < y(v^*);$$

i.e., $y(v^*) - y^-(v^*) = P\{U_1 + v^* = U_2\} > 0$. Let $p = (r_1 - y^-(v^*))/(y(v^*) - y^-(v^*))$. Note that $0 \leq p \leq 1$. The opportunistic scheduling policy is then given by

$$Q^*(\vec{U}) = \begin{cases} 1 & \text{if } U_1 + v^* > U_2 \\ 1 & \text{with prob. } p \text{ if } U_1 + v^* = U_2 \\ 2 & \text{with prob. } 1 - p \text{ if } U_1 + v^* = U_2 \\ 2 & \text{if } U_1 + v^* < U_2 \end{cases}$$

It is clear that the policy $Q(\vec{U})$ defined above is feasible:

$$P\{Q(\vec{U}) = 1\} = P\{U_1 + v^* > U_2\} + P\{U_1 + v^* = U_2\}p = r_1.$$

The policy can be described as follows. The space spanned by U_1 and U_2 is divided into two halves by the line $U_1 + v^* = U_2$. Above the line (i.e., $U_2 > U_1 + v^*$), we always schedule user 2 to transmit. Under the line (i.e., $U_1 + v^* > U_2$), we always schedule user 1 to transmit. If the probability of the line is positive, some randomization is needed if we fall on the line—with probability p, we schedule user 1 and with probability $1 - p$, we schedule user 2, where $p = (r_1 - y^-(v^*))/(y(v^*) - y^-(v^*))$ is determined by the temporal fairness constraints.

3.2.2 General Case

Now we extend the policy from the previous section to the N-user case. Define

$$y_i(\vec{v}) = P\{U_i + v_i \geq \max_{j \neq i}(U_j + v_j)\}, \quad \text{for } i = 1, \cdots, N,$$

where $\vec{v} = (v_1, \cdots, v_N)$. Note that $y_i(\vec{v})$ is a monotonically-increasing right-continuous function of v_i and a monotonically-decreasing left-continuous function of v_j, $j \neq i$.

In Appendix A.5, we show that there exists a $\vec{v}^*$ that satisfies $P\{Q^*(\vec{U}) = i\} = r_i$, where the opportunistic policy is

$$Q^*(\vec{U}) = \underset{i}{\operatorname{argmax}}(U_i + v_i^*). \tag{3.2}$$

In the argmax above, we break ties probabilistically by picking a user i among those that achieve the maximum above with a certain probability.

Note that $\vec{v}^*$ is not unique. There are N components but only $N-1$ independent constraint equations: $P\{Q(\vec{U}) = i\} = r_i$, for $i = 1, \cdots, N-1$, and $P\{Q(\vec{U}) = N\} = 1 - \sum_{i=1}^{N-1} r_i$ is a linear combination of the first $N-1$ equations, we can simply set $v_N = 0$.

The policy Q^* defined in (3.2), which represents our opportunistic scheduling policy, is optimal in the following sense.

Proposition 1 *The policy Q^* is a solution to the problem defined in (3.1); i.e., it maximizes the average system performance under the temporal fairness constraint.*

For a proof of the above proposition, see Appendix A.1. The parameter $\vec{v}^*$ is the "offset" used to satisfy the temporal constraint. Under this constraint, the scheduling policy schedules the "relatively-best" user to transmit. User i is the "relatively-best" user if $U_i + v_i^* \geq U_j + v_j^*$ for all j. In a special case where $v_j^* = 0$ for all j, the scheduling policy reduces to $Q(\vec{U}) = \operatorname{argmax}_i U_i$; i.e., always schedules the user with the largest performance value to transmit.

Proposition 2 *If the performance values of different users are independent, then*

$$E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right) \geq r_i E(U_i),$$

where $E\left(\mathbf{1}_{\{Q(\vec{U})=i\}}\right) = r_i$.

For a proof of this proposition, see Appendix A.2. Note that $E(U_i \mathbf{1}_{\{Q(\vec{U})=i\}})$ is the average performance of user i of the opportunistic scheduling policy and $r_i E(U_i)$ is the performance of a non-opportunistic scheduling scheme (e.g., round-robin). This proposition studies the individual performance of each user. If users' performance values are independent of each other, every user will get the same or better performance in the opportunistic scheduling scheme than in round-robin.

3.2.3 Parameter Estimation

Basically, the opportunistic scheduling policy is given by

$$Q^*(\vec{U}) = \underset{i}{\operatorname{argmax}}(U_i + v_i^*),$$

where the v_i^*s are parameters determined by the distribution of $\vec{U}$. In practice, this distribution is unknown, and hence we need to estimate the parameters v_i^*, for $i = 1, \cdots, N-1$. Figure 3.1 shows a block diagram of a practical scheduling procedure that incorporates on-line estimation of these parameters.

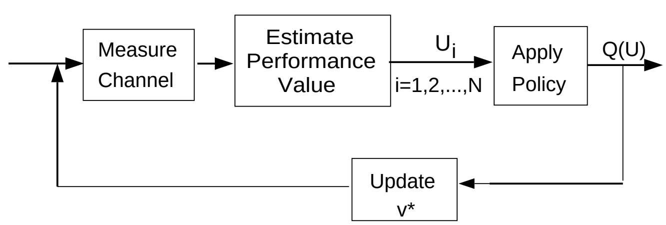

Fig. 3.1. Block diagram of the scheduling policy with on-line parameter estimation

In this section, we focus on the block that implements the on-line estimation of the parameters v_i^*, $i = 1, \cdots, N-1$, labeled "Update v^*" in Figure 3.1. We use a standard stochastic approximation algorithm to estimate $\vec{v}^*$ (the vector of the v_i^*).

We first explain intuitively the idea of the stochastic approximation algorithm used here. For a systematic and rigorous study of stochastic approximation algorithms, see [62, 63]. Suppose we want to solve the root-finding problem $f(x^*) = 0$, where f is a continuous function with one root x^* (both x^* and $f(x^*)$ are vectors of the same dimension). If we can evaluate the value of $f(x)$ at any x, then we can use the iterative algorithm

$$x^{k+1} = x^k - a^k f(x^k),$$

which will converge to x^* as long as the step size a^k is appropriately chosen; e.g., $a^k = 1/k$. Suppose that we cannot obtain exactly the value of $f(x^k)$ at x^k, but instead we only have a noisy observation g^k of $f(x^k)$ at x^k; i.e., $g^k = f(x^k) + e^k$ where e^k is the observation error (noise). In this case, it is well-known that if $E(e^k) = 0$ (i.e., the mean of the observation error is zero), then the algorithm

$$x^{k+1} = x^k - a^k g^k,$$

converges to x^* with probability 1 under appropriate conditions on a^k and f (see, e.g., [62, 63]).

We use a stochastic approximation algorithm to estimate $\vec{v}^*$. For this, note that we can write $\vec{v}^*$ as a root of the equation $f(\vec{v}^*) = 0$, where the ith component of $f(\vec{v}^*)$ is given by

$$f_i(\vec{v}^*) = P\{Q^*(\vec{U}) = i\} - r_i, \qquad i = 1, \cdots, N-1$$

and

$$Q^*(\vec{U}) = \underset{i}{\operatorname{argmax}}(U_i + v_i^*).$$

We use a stochastic approximation algorithm to generate a sequence of iterates $\vec{v}^1, \vec{v}^2, \cdots$ that represent estimates of $\vec{v}^*$. Each $\vec{v}$ defines a policy Q^k given by

$$Q^k(\vec{U}) = \underset{i}{\operatorname{argmax}}(U_i + v_i^k).$$

To construct the stochastic approximation algorithm, we need an estimate g^k of $f(\vec{v}^k)$. Note that although we cannot obtain $f(\vec{v}^k)$ directly, we have a noisy observation of its components:

$$g_i^k = \mathbf{1}_{\{Q^k(\vec{U})=i\}} - r_i, \qquad i = 1, \cdots, N-1.$$

The observation error in this case is

$$e_i^k = g_i^k - f_i(\vec{v}^k) = \mathbf{1}_{\{Q^k(\vec{U})=i\}} - P\{Q^k(\vec{U}) = i\},$$

and thus we have $E(e_i^k) = 0$. Hence, we can use a stochastic approximation algorithm of the form

$$v_i^{k+1} = v_i^k - a^k \left(\mathbf{1}_{\{Q^k(\vec{U})=i\}} - r_i \right),$$

where $a^k = 1/k$. For the initial condition, we can set v_i^1 to be 0, or some estimate based on the measurement history. For the above algorithm, following the standard proof of [62], we can show that $\{v_i^k\}$ converges to v_i^* with probability 1. Furthermore, to accelerate the convergence and to reduce the range of the fluctuation of the stochastic approximation algorithm, we can use the standard technique of averaging (see, e.g., [63]):

$$\bar{v_i}^k = \left(1 - \frac{1}{k}\right) \bar{v_i}^{k-1} + \frac{1}{k} v_i^k.$$

Simulations show that with the stochastic approximation algorithm, v_i^k converges to $\vec{v}^*$ relatively quickly.

When U_is are not continuous random variables, there may be a "tie" case. Specifically, $P\{U_i + v_i^* = \max_{j \neq i}(U_j + v_j^*)\} > 0$, for some i, we break ties probabilistically by picking a user i among those that achieve the maximum above with a certain probability. Theoretically, v_i^k will converge to v_i^* and we should estimate the tie-break probability. In practice, this "tie-break" can be handled by the adaptive nature of stochastic approximation. In this case, v_i^k fluctuates around v_i^* within a small range. The idea is: when v_i^k is getting larger, $P\{Q = i\} > r_i$, hence, v_i^k will be dragged down and so on. Simulations show that stochastic approximation works well in the case with "tie-break".

As we will see in Section 3.5, simulations results show that the system performance obtained in our simulation is very close to that of the true optimal value, which implies that stochastic approximation works well in our situation.

3.3 Implementation Considerations

So far, we have described our scheduling policy, proved its optimality (in the appendix), and addressed the problem of estimating the parameter values needed for the policy. In this section, we explore some implementation considerations for our scheduling policy.

In our scheduling policy, the base station needs to obtain information of each user's performance value at a given time-slot to make the scheduling decision. The performance value of a user can be estimated either by the user or by the base station, based on the channel condition and/or measurements from previous transmissions. For the downlink case, a user could measure the received signal power level (from the user's base station) and the interference power level. The user could then calculate the performance value of the time-slot based on the channel condition and other factors (such as power consumption). For example, suppose a user's performance is defined as its throughput, which is an S-shape function of the SINR, as shown in Figure 2.1. Based on the estimated SINR, the user can then obtain its performance value. For the uplink case, the base station could estimate the user's channel condition based on the received signal from the user. Assuming the base station knows the form of the performance value for each user (i.e., how the performance value depends on the SINR and/or other factors), the performance value could then be calculated by the base station.

If the performance value is estimated by the user, this information needs to be sent to the base station, which can be accomplished in several ways. For example, each user could maintain a small signaling channel with the base station. Alternatively, the required information could be piggybacked over the user's acknowledgment packets.

As mentioned before, the length of a time-slot in our scheduling policy can be different from an actual time-slot of the physical channel. The length of a scheduling time-slot depends on how fast the channel condition varies and how fast we want to track the variation. The usual tradeoff between accuracy and signaling overhead exists here. Specifically, more frequent updating provides more accurate tracking of varying channel conditions, but incurs higher signaling costs. In practice, to decrease signaling costs, a user can update its information only when the change in the performance value is larger than a certain threshold. Furthermore, it is not necessary for all users to update at the same time. Note that we ignore propagation delay in the dissertation.

In the following we summarize our scheduling procedure, which incorporates the on-line parameter estimation algorithm described in the last section. As mentioned before, the initial value of $\vec{v}^*$ can be set to $\vec{0}$ or some estimate based on history information. At each time-slot $k = 1, 2, \cdots$, the system performs the following steps:

1. Estimate U_i^k;
   - Uplink: the base station estimates each user's channel condition and calculates the values of U_i^k, $i = 1, \cdots, N$;
   - Downlink: user i measures its channel condition, calculates U_i^k, and informs the base station;
2. The base station decides which user should be scheduled to transmit in the time-slot based on the scheduling policy:
$$Q^k(\vec{U}^k) = \underset{i}{\operatorname{argmax}}(U_i^k + v_i^k);$$
3. The scheduled transmission takes place;
   - Uplink: the base station broadcasts the ID of the selected user and the selected user transmits in the time-slot.
   - Downlink: the base station transmits to the selected user;
4. The base station updates the parameter vector $\vec{v}^{k+1}$ via
$$v_i^{k+1} = v_i^k - a^k \left(\mathbf{1}_{\{Q^k(\vec{U}^k)=i\}} - r_i \right);$$
   For the stationary case, we set $a^k = 1/k$. For the nonstationary case, we set a^k to a small constant to track system variations.

Note that the computation burden above is $O(N)$ per time-slot, where N is the number of users sharing the channel (usually on the order of tens), which suggests that the procedure is easy to implement in practice.

3.4 Time-Fraction Assignment

Recall that our opportunistic scheduling scheme assumes a given temporal fairness requirement. The temporal fairness constraint $r_1, \ldots, r_N$ represents a prespecified time-fraction assignment. In the following, we describe three time-fraction assignment schemes, which approach the problem from different viewpoints.

User bidding: Suppose m_i is the amount of money that user i is willing to pay per unit time to access the wireless resources. The network then assigns time-fractions to users in proportion to their willingness to pay:

$$r_i = \frac{m_i}{\sum_{j=1}^{N} m_j}.$$

Hence, the more a user is willing to pay, the higher the fraction of the resources assigned to the user.

Fair sharing: If there are N users in the system, each user is assigned $r_i = 1/N$; i.e., each user receives the same share of resources. This scheme provides fair resource sharing assuming users are homogeneous. We can extend this scheme to the multi-class case as follows. Suppose there are L classes of users, where each class has l_i active users, and an associated weight w_i reflecting the importance and/or resource requirement of this class. Then the time-fraction assignment for a user in class i is:

$$r_i = \frac{w_i}{\sum_{j=1}^{L} w_j l_j}.$$

Bias sharing: The philosophy here is to allocate resources to users in proportion to their expected performance values. The corresponding time-fraction assignment is given by

$$r_i = \frac{E(U_i)}{\sum_{j=1}^{N} E(U_j)}.$$

This scheme clearly favors users with high performance values, but at the same time does not totally ignore the requirement of users with poor performance values.

3.5 Simulation Results

In this section, we present numerical results from computer simulations of our scheduling scheme. The distinctive feature of our scheduling policy is that it exploits time-varying channel conditions—the policy dynamically decides which user should be scheduled to transmit in a time-slot based on users' current performance values. For the purpose of simulations, we assume that the time-fraction assignment is done using *fair sharing*, i.e., the total resources are evenly divided among the users. Note that the policy that shares the resource (time in this case) in this manner, but does not exploit channel conditions is the well known *round-robin* scheme. Hence, to evaluate the performance gain of our dynamic and opportunistic assignment of transmissions, we compare the performance of our policy with that of the round-robin scheme. We will show three sets of simulation results. The first one is to simulate an actual cell and to show how much improvement our scheduling scheme will have. We then show how estimation errors affect the scheduling scheme and how stochastic approximation works in a continuous case. The last simulation is to show how stochastic approximation works in a "tie-break case".

3.5.1 An Actual Cell

Our simulation environment is described in the following. We consider a multi-cell system consisting of a center cell surrounded by hexagonal cells of the same size. The base station is at the center of each cell and simple omni-directional antennas are used by mobiles and base stations. We focus on the performance of the downlink of the center cell because downlink communication is more important for data services. The frequency reuse factor is 3 and co-channel interference from the first-tier neighboring cells is taken into account. We assume that each cell has a fixed number of frequency bands. Usually there are tens of users in each cell sharing different frequency bands. We focus on one frequency band, which is shared by 25 users in the central cell. The scheduling policy decides which user should transmit in this frequency band at each time-slot. The users have exponentially distributed "on" and "off" periods.

We model user mobility as follows. The velocities of mobile users are independent random variables uniformly distributed between the minimum (2km/h) and the maximum velocity (100km/h). The directions of mobile users are independent random variables uniformly distributed between 0 and 2π. A mobile user chooses its velocity when it becomes active and the velocity is fixed during that on-period. The direction of a mobile user changes periodically. When a user becomes active, its location is uniformly distributed in the cell. If a user moves out of the border, we assume that it reappears at a point that is symmetric to the exiting point about the center base station.

As mentioned earlier, for our simulation experiments, we use the *fair sharing* time-fraction assignment scheme. When the number of active users changes, i.e., when an active user becomes nonactive or vice versa, we update the temporal fairness requirement r_i for all active users. In other words, if N is the number of active users sharing the channel in the central cell, we set $r_i = 1/N$, for all i.

The channel gains of the users are mutually independent random processes determined by the sum of two terms: one due to path (distance) loss and the other to shadowing. To be conservative, we assume, in the simulation, that the effects of fast multipath fading are averaged out via interleaving, diversity, etc., because current standards have not specified faster than frame-rate adaptation. However, we should note that fast fading is considered in the Qualcomm/HDR proposal [25], and if fast fading could be tracked, our scheme could provide even higher performance improvements than shown here.

We adopt the path-loss model (Lee's model) and the slow log-normal shadowing model in [32], as discussed in Section 2.1. Specifically, the channel gain $g(k)$ (in dB) at time-slot k between an arbitrary user at a distance d from a base station is given by:

$$g(k) = l_p(k) + s(k).$$

where $l_p(k)$ and $s(k)$ are terms representing path-loss and shadowing, respectively. The path loss $l_p(k)$ (dB) is obtained as

$$l_p(k) = K + 38.4 \log_{10}(d(k)) - \alpha_0,$$

where α_0 is a correction factor used to account for different base station and mobile station (MS) antenna heights, transmit powers, and antenna gains, and $K = 103.41$ is a constant in the simulation assuming that the transmission power of a base station is fixed at 10W.

Shadowing is the result of the transmitted signal passing through or reflecting off some random number of objects such as buildings, hills, and trees. The shadowing term $s(k)$ (in dB) is usually modeled as a zero-mean stationary Gaussian process with autocorrelation function given by

$$E(s(k)s(k+m)) = \sigma_o^2 \xi_d^{vT/D},$$

where ξ_d is the correlation between two points separated by a spatial distance D (meters), and v is velocity of the mobile user. In our simulation, we use a value of $\sigma_o = 4.3$ dB, corresponding to a correlation of 0.3 at a distance of 10 meters, as reported by Gudmundson [34].

The parameters of the simulation and their values are summarized in Table 3.1.

In the following, we describe the simulation procedure in detail. At the beginning of the simulation, we set $\vec{v}^1 = \vec{0}$. We maintain an ordered list of users in the system. Let N be the number of active users. Each active user has a temporal requirement of $1/N$. At each time-slot $k = 1, 2, \cdots$, the following steps are simulated:

1. If user i is active, we generate U_i^k. In our simulation, a user's performance is a function of SINR, as shown in Figure 3.2. Each user measures the received power level from the central base station, and the interference power level received from neighboring cells. Based on these measurements, the user calculates the SINR, and thus the corresponding performance value as a function of SINR. Figure 3.2 shows the forms of the performance values used by different users.

Cell radius	1000m
Propagation environments	North American suburb
Frequency	1845 MHz
Distance Exponent (β)	3.84
Height of base station antenna	38.4m
Height of MS antenna	1.5m
base station transmission power	10W
base station antenna gain	6dB above dipole gain
MS antenna gain	0dB above dipole gain
Shadowing standard deviation	4.3dB
Shadowing correlation distance	10m
Background noise power	120dBm
Minimum speed	2km/h
Maximum speed	100km/h
Period for direction change	127 steps
Sample step size	10ms
Mean duration of on-period	5000 steps (50s)
Mean duration of off-period	2500 steps (25s)
Stochastic approximation step	0.01

Table 3.1
Simulation parameters and values

To avoid crowding the figure, we only show the performance functions of 8 users. The performance values of users 1 and 2 are step-functions of their SINR, and user 2 has a higher threshold than user 1. The performance values of users 3–4 are linear functions of their SINR (in dB), with different slopes. Users 5–8 have performance values that are S-shape functions of their SINR, with different parameters. Totally, there are 4 users with step-functions, 6 users with linear functions, and 15 users with S-shape functions in the simulation.

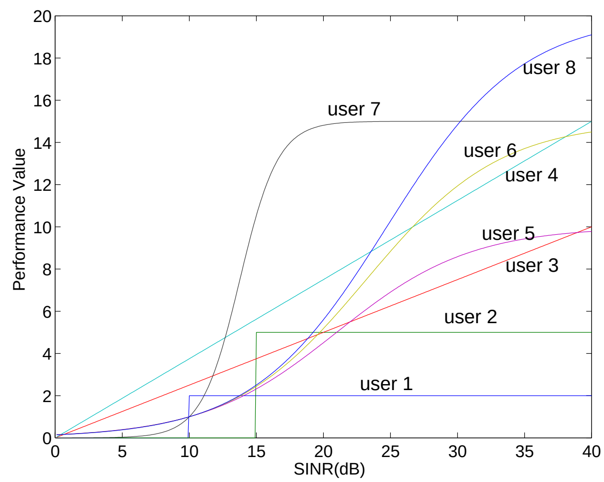

Fig. 3.2. Users' performance values as a function of SINR.

2. Active users transmit their values of U_i^k to the base station through a signaling channel.

3. Based on the vector of performance values $\vec{U}^k$, the base station decides which user to schedule in the time-slot:

$$Q^k(\vec{U}^k) = \underset{i \in A}{\operatorname{argmax}}(U_i^k + v_i^k),$$

where A is the index set of all active users.

4. If user $j = Q(\vec{U}^k)$ is the selected user, then the base station transmits to user j in the time-slot k. The system receives a performance "reward" equal to the performance value U_j^k.

5. In the round-robin scheduling scheme, we set J to be the index of the next active user in our ordered list of users, and let user J transmit. The system receives a performance "reward" equal to the performance value U_J^k of user J.

6. The base station updates $\vec{v}^{k+1}$ for all active users as follows:

$$v_i^{k+1} = v_i^k - a^k \left(\mathbf{1}_{\{Q^k(\vec{U}^k)=i\}} - r_i \right).$$

   Because we are simulating a non-stationary system, we set $a^k = 0.01$ to track changes in the system. In general, the larger the value of a^k, the faster that v_i^k tracks v_i^*, but at the same time the larger the fluctuation of v_i^k around the value of v_i^*.

The system performs the above procedure for every time-slot. Whenever the number of active users changes, the base station updates the temporal fairness requirements according to the fair sharing scheme, and the value of $\vec{v}^k$ at that time is used as the initial value of the on-line parameter estimation procedure in the new system state.

Figure 3.3 shows the results of our simulation experiment. In the figure, the x-axis represents the users' IDs. For each user, we compare the average performance in our opportunistic scheduling policy (the first bar) with that of the round-robin (RR) policy (the second bar). We can see that in every case, our opportunistic policy significantly outperforms the round-robin policy, with gains of 20% to 150%. The amount of improvement varies from user to user because different users have different performance functions. The third bar in the figure is the ratio of the total number of slots assigned to each user in our opportunistic scheme to that of the round-robin scheme, which was set equal to that required by the temporal fairness constraint. For all users, the third bar is virtually identical to 1. Hence, our scheduling scheme satisfies the temporal fairness constraint, which indicates that our stochastic approximation algorithm works well in the simulation experiment even in the nonstationary case.

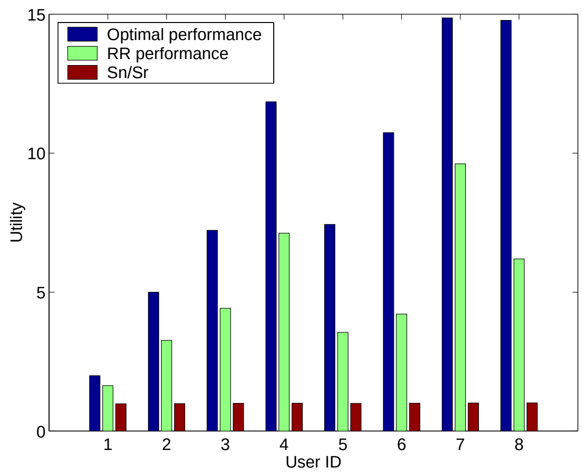

Fig. 3.3. Comparison of the opportunistic scheduling policy with the round-robin scheme. In the figure, S_n is the number of time-slots assigned to user i in the optimal scheduling policy and S_r is the number of slots assigned to user i in the round-robin scheme.

3.5.2 Estimation Error

We next show results for an experiment designed to evaluate the impact of our online parameter estimation procedure and the sensitivity of our opportunistic scheduling scheme to estimation errors. There are two parts of estimation errors: the parameter estimation error (i.e., the difference between $\vec{v}^*$ and its estimate $\vec{v}^k$) and the performance-value estimation error (i.e., the difference between U_i and its estimate $\hat{U}_i$).

ID	mean	autoreg. coefficient	stand. deviation	σ_n^i
1	10	0.3	10.8	20
2	10	0.4	6.9	16
3	10	0.5	4.0	12
4	10	0.6	2.5	10

Table 3.2
Gaussian Process Parameters

We generate four time-correlated Gaussian processes, representing the performance-value sequences for four users. Let η_i be the auto-regression correlation factor of user i. We have

$$U_i^{k+1} = \eta_i U_i^k + (1 - \eta_i) n_i^k,$$

where $\{n_i^k\}$ is a sequence of i.i.d. Gaussian random variables with mean value of 10 and standard deviation σ_n^i, as shown in Table 3.2. We also display in Table 3.2 the means and standard deviations of the Gaussian processes for the four users. Each user has exponentially distributed "on" and "off" periods, with mean value 5000 and 2500 time-slots, respectively.

We compare the results for four different cases.

- Ideal case; i.e., Q^* with known threshold $\vec{v}^*$ assuming exact values of U_i are known;
- Estimated thresholds assuming exact values of U_i are known;
- Estimated thresholds with estimated value of U_i; i.e., $\hat{U}_i = U_i + e_i$, where e_i is the estimation error, $e_i \sim N(0, 4)$, and e_is are independent;
- Estimated thresholds with estimated value of U_i; i.e., $\hat{U}_i = U_i + e_i$, where e_i is the estimation error, $e_i \sim N(0, 8)$, and e_is are independent.

Figure 3.4 shows the average performance from the above four cases. The first bar is the normalized performance value under the ideal condition; the second bar is that of estimated thresholds with ideal measurements (i.e., the exact value of U_i is known); the third and fourth bars represent the normalized performance with different estimation errors; i.e., $\hat{U}_i = U_i + e_i$, where $e_i \sim N(0,4)$ and $e_i \sim N(0,8)$ respectively. We can see in Figure 3.4 that the performance of both the optimal policy Q^* and our policy Q^k with estimated parameters are quite comparable, and are significantly higher than that of the round-robin policy. This suggests that our on-line parameter estimation scheme works well and that the parameter estimation errors present do not significantly degrade the performance of the policy relative to the optimal policy.

Moreover, the performance gains appear to be related to the standard deviation: the higher the standard deviation, the larger the performance gain. Note that in the round-robin scheme, the performance levels for all users are all approximately equal to the mean of the Gaussian processes (which is the mean performance value). This is to be expected because the round-robin scheme allocates an equal fraction of time-slots to each user, regardless of the channel conditions. Our opportunistic approach takes advantage of favorable transmission conditions, thereby leading to average performance values that are far above the mean of the Gaussian processes. The third bar shows the average performance of users with estimation error $e_i \sim N(0,4)$. With this estimation error, the average performance is still close to that of the optimal case. When the estimation error increases, it is not surprising that the average performance decreases. The fourth bar shows a situation with very large estimation errors, especially for user 4. However, even in this case, our scheduling scheme still outperforms that of round-robin. Hence, the opportunistic scheduling scheme is robust to estimation errors.

Figure 3.5 shows the ratio of the time-fractions obtained by our policy Q^k to that of the round-robin policy (which, as pointed out before, are equal to the prespecified values). As we can see, our scheme satisfies the temporal fairness constraints very well, even in the case with very large estimation errors.

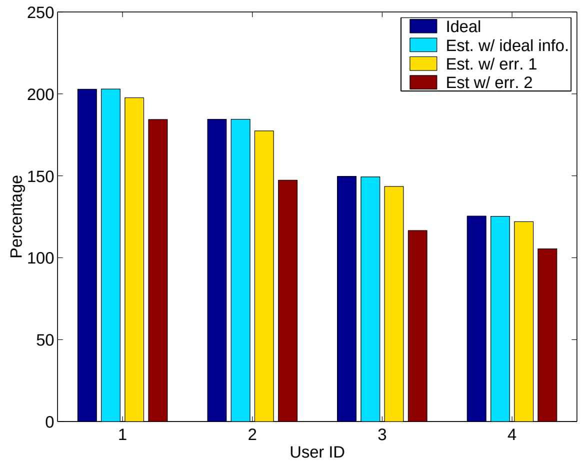

Fig. 3.4. Average performance value, normalized over round-robin.

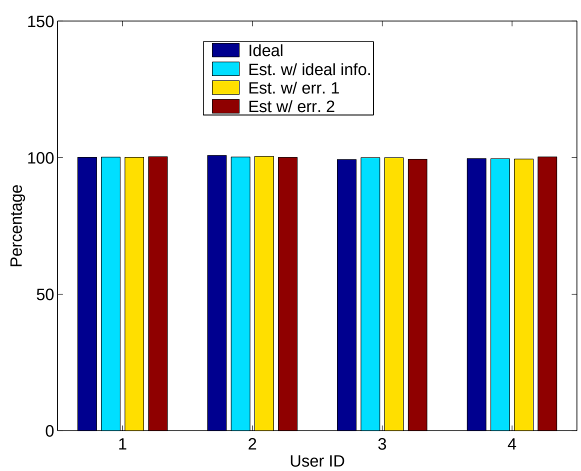

Fig. 3.5. Fairness, normalized over round-robin

Lastly, Figure 3.6 shows the estimated threshold of a user when the system status changes (i.e., some user becomes active or some user becomes inactive).

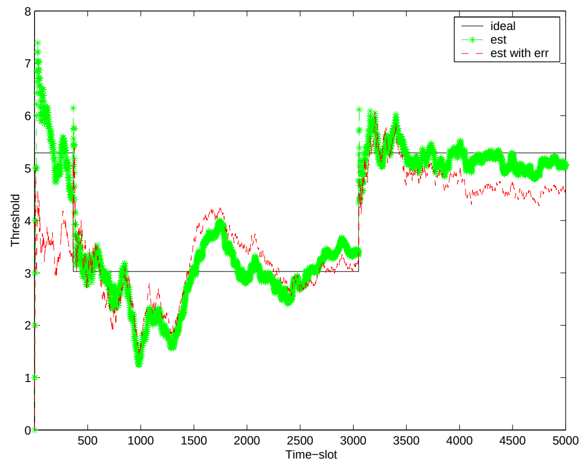

Fig. 3.6. Estimated threshold when the number of users changes.

3.5.3 Tie-break Case

In this simulation, we generate four discrete random variables to represent the performance values of four users. The distribution functions of these four random variables are shown in the following.

$$U_1 = \begin{cases} 1 & \text{with prob. } 0.5 \\ 2 & \text{with prob. } 0.5 \end{cases} \quad U_2 = \begin{cases} 1 & \text{with prob. } 0.5 \\ 3 & \text{with prob. } 0.5 \end{cases}$$
$$U_3 = \begin{cases} 2 & \text{with prob. } 0.5 \\ 3 & \text{with prob. } 0.5 \end{cases} \quad U_4 = \begin{cases} 2 & \text{with prob. } 0.5 \\ 4 & \text{with prob. } 0.5 \end{cases},$$

where U_is are independent. It is trivial to show that $\vec{v}^* = [2, 1, 1, 0]$.

A tie-break case occurs. For example, when $U_1 = 1, U_2 = 3, U_3 = 2, U_4 = 4$, $U_2 + v_2^* = U_4 + v_4^* = \max(U_i + v_i^*)$. In this case, with probability 0.5, we should let user 2 transmit, and with probability 0.5, we should let user 4 transmit. In general, $P\{U_i + v_i^* = \max_{j \neq i}(U_j + v_j^*)\} > 0$, for some i, we break ties probabilistically

ID	T_{RR}	T_{opt}	T_{est}	R_f
1	1.5	1.875	1.874	1.002
2	2	3	3.000	1.004
3	2.5	2.875	2.870	1.002
4	3	4	4.000	0.992

Table 3.3
Comparison of the average performance of scheduling policy Q^* (with known $\vec{v}^*$), Q^k (with estimated $\vec{v}^k$), and round-robin.

by picking a user i among those that achieve the maximum above with a certain probability. Theoretically, v_i^k will converge to v_i^* and we should estimate the tie-break probability. In practice, this "tie-break" can be handled by the adaptive nature of stochastic approximation. In this case, v_i^k fluctuates around v_i^* within a small range. The idea is: when v_i^k gets larger, $P\{Q = i\} > r_i$, hence, v_i^k will be dragged down and so on. This simulation is designed to test the stochastic approximation algorithm in a tie-break case.

The simulation is run for 2,000 time-slots with initial value $\vec{v}^* = 0$. Table 3.3 shows the average performance value and fairness obtained from the simulation. We simulate our scheduling procedure where we use our on-line parameter-estimation algorithm to estimate $\vec{v}^*$ based on measurements. The average performance obtained is shown as T_{est} in Table 3.3 and the optimal performance (with known $\vec{v}^*$ and know tie-break probabilities) is noted as T_{opt}. As a benchmark comparison, we also simulated the average performance of the round-robin policy, represented by T_{RR} Table 3.3. In the table, R_f is the ratio of the number of time-slots the user gets in the simulation over the number of time-slots the user should get according to its fair share. We can see T_{opt} and T_{est} are close and R_f is close to one. Furthermore, in Figure 3.7, we can see that v_i^k converges to v_i^* fairly fast. Hence, stochastic approximation works well in this case.

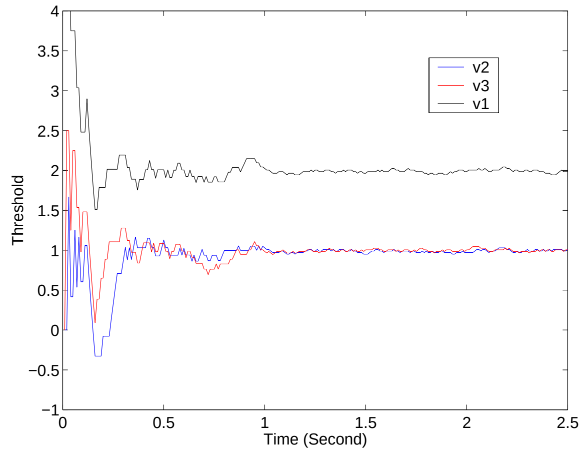

Fig. 3.7. Convergence of the threshold values in a discrete case

In this section, we present simulation results of our opportunistic scheduling scheme under various conditions. It shows that our scheduling scheme works well in the case where stationarity is not guaranteed, and improves system performance significantly, even in the presence of estimation errors.

3.6 Short-term Performance

The scheduling scheme described thus far meets the long-term performance requirements of users; i.e., the long-term average of the fraction of time-slots assigned to a user is guaranteed. However, with such a scheme, it is possible that a user could be starved for a long time (say, more than a few seconds), which may be undesirable for certain users. Usually, a user may also have the demand for good “short-term” performance — the user expects that the amount of service obtained within a relatively short time interval be close to it amount it should get.

In the GPS (Generalized Processor Sharing) model [29], each flow is treated as a fluid flow. Each flow i is given a weight w_i, and for any time interval $[t_1, t_2]$ during which both sessions i and j are continuously backlogged, the resource granted to each flow i, $G_i(t_1, t_2)$, satisfies the following property:

$$\left|\frac{G_i(t_1,t_2)}{w_i} - \frac{G_j(t_1,t_2)}{w_j}\right| = 0. \tag{3.3}$$

This means that a user gets its fair share of resource during any time interval. There is an alternative to (3.3). Let t_0 be a starting point such that from time t_0 onwards, both sessions i and j are continuously backlogged. It is clear that the satisfaction of (3.3) is equivalent to:

$$\left|\frac{G_i(t_0,t)}{w_i} - \frac{G_j(t_0,t)}{w_j}\right| = 0 \tag{3.4}$$

for all $t > t_0$, and users i and j are both continuously backlogged during the time interval $[t_0, t]$.

We extend the above concept to a time-slotted system, where one time-slot is exclusively used by one user. Let w_i be the weight of user i and $r_i(k)$ be the temporal fairness requirement of user i at time k. Note that when the set of active users change, $r_i(k)$ may change. For example, following the tradition in GPS, $r_i(k)$ can be set as:

$$r_i(k) = \frac{w_i}{\sum_{j \in A_k} w_j},$$

where A_k is the set of active users at time k. User i is guaranteed a minimum share of the resource $w_i / \sum_{j \in A} w_j$ during its active period, where A is the set of all users.

Let k_i be the time that user i becomes active. Suppose that during the time interval $[k_0, k]$, both users i and j are continuously active, where $k_0 = \max(k_i, k_j)$. Let $S_i(k_0, k)$ be the number of time-slots assigned to user i from k_0 to k. An approximation of (3.4) is:

$$\left|\frac{S_i(k_0,k)}{w_i} - \frac{S_j(k_0,k)}{w_j}\right| \leq \Gamma,$$

where $\Gamma \geq 0$ is a constant.

Let $F_i(k_i, k)$ be the counter of the resource entitled of user i; i.e., the number of time-slots that should be assigned to user i during the time interval $[k_i, k]$:

$$F_i(k_i,k) = \sum_{t=k_i}^{k} r_i(t) = \sum_{t=k_i}^{k} \frac{w_i}{\sum_{j \in A_t} w_j}, \tag{3.5}$$

where A_t is the set of active users at time t. It is obvious that the $F_i(k_i, k)$ satisfy (3.4) and hence (3.3) at each discrete time k. Note that $F_i(k_i, k)$ may not be an integer, and thus may not be achievable when a time-slot is exclusively used by one user.

We use $F_i(k_i, k)$ as a benchmark. To improve the short-term performance, we want $S_i(k_i, k)$ to be close to $F_i(k_i, k)$. We modify our previous opportunistic scheduling scheme in the following way. Let

$$\Delta_i^k = F_i(k_i, k) - S_i(k_i, k).$$

If $\Delta_i^k > 0$, then user i is "lagging" (i.e., the user gets less resource than it should get), and if $\Delta_i^k < 0$, then user i is "leading". The idea is to increase the probability of transmission of a lagging user and decrease the probability of transmission of a leading user. Hence, a direct modification of our scheduling policy is the policy B^k given by

$$B^k(\vec{U}^k) = \underset{i=1,\cdots,N}{\operatorname{argmax}} (U_i^k + v_i^*) \left(\frac{\Delta_i^k}{\alpha} + \beta \right), \tag{3.6}$$

where $\min v_i^* = 0$, and α and β are positive constants. When the value of α is smaller, the effect of Δ_i^k is more significant, and thus the short-term performance is better. The value of β acts as a threshold — a user is forbidden to transmit if the amount by which it leads is greater than $\beta\alpha$.

We next consider the case where there are changes in the set of active users. When a new user comes into the system, the system adjusts the r_i for all users, and the new user j starts a counter $F_j(k_j, k)$ for its resource share. When a user leaves the system, the system adjusts the r_i and the counters of fair share for other active users. Suppose user i leaves the system at time k and user i has been served with $S_i(k_i, k)$ time-slots. Recall that $\Delta_i^k = F_i(k_i, k) - S_i(k_i, k)$; i.e., user i has Δ_i^k time-slots less than its share. Because the user is gone, we cannot enforce $S_i(k_i, k)$ to be close to $F_i(k_i, k)$ anymore. This discrepancy has to be absorbed by other active users. We update the counter $F_j(k_j, k)$ of any active user j by replacing the value of $F_j(k_j, k)$

by $F_j(k_j, k) + \Delta_i^k/N_{act}$, where N_{act} is the number of active users. In other words, the discrepancy is evenly distributed among all active users. Note that this way of handling users' departures is intuitive, but not necessarily optimal. Actually, it is challenging to even define a good optimal criterion in the situation where there exists the tradeoff between short-term performance and the overall system performance.

We use the same simulation setup as in 3.5.2. Four Gaussian random processes are used to represent the performance-value sequences of four users, and their parameters are shown in Table 3.2. The simulation runs for 1,000,000 time-slots (while all four users have exponentially distributed on-offs). Next, we show two metrics for the short-term performance for user 4, which has an auto-regression coefficient of 0.6. Note that user 4 has the worst short-term performance among all four users because it has the highest correlation across time.

The first metric is the starving-time, defined as the time interval between two contiguous time-slot assignments when the user is active. Note that starving-time is closely related to the delay a user experiences. Figure 3.8 shows the starving-time histogram. In the legend, MOS represents the modified opportunistic scheduler defined in (3.6); OSI is the ideal opportunistic scheduler with known threshold $\vec{v}^*$; OSE is the opportunistic scheduler using stochastic approximation to estimate the threshold; and IND represents the numerical result when a user's performance values at different time-slots are independent. If a user's performance values are independent across time, the starving-time is binomially distributed; i.e., $P\{\text{starving-time} = n\} = (1 - r_i)^{(n-1)} r_i$, where r_i is the temporal fairness requirement of user i. Because user 4's performance value is correlated across time, compared with the IND case, the probability of a long starving-time of user 4 in the OSI and OSE cases, which do not consider the short-term performance, is larger, but it is also dropping exponentially. Furthermore, the probability that a large starving-time occurs in MOS is much smaller than that of OSI and OSE, and it is also smaller than that of IND. Hence, the chance that a user is starved decreases and the short-term performance is improved.

The second metric of the short-term performance is defined as

$$\frac{E[(S_i(k,k+\Delta m) - F_i(k,k+\Delta m))^2]}{\Delta m},$$

where Δm is the length of the window by which we measure the discrepancy between the fair share $F_i(k,k+\Delta m)$ and $S_i(k,k+\Delta m)$ while user i is active during the interval $[k,k+\Delta m]$. Because $E[S_i(k,k+\Delta m)] = F_i(k,k+\Delta m)$, this metric is the variance of $S_i(k,k+\Delta m)$ normalized over the window size Δm. The discrepancy should be zero in the Fluid Fair Model and is no larger than 1 in the round-robin scheduling scheme. In Figure 3.9, we show that MOS results in noticeable decrease of the normalized variance.

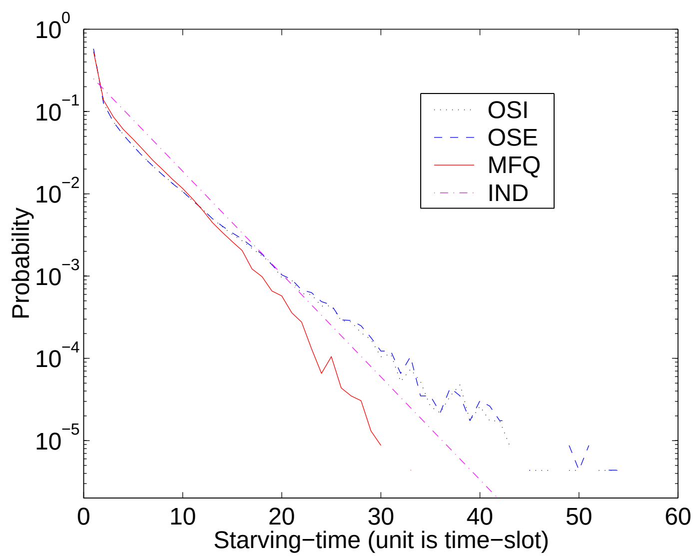

Fig. 3.8. Starving-time histogram

We should mention here that the modified scheduling scheme does not decrease the average performance significantly. In this simulation, OSI outperforms round-robin by 25% in terms of the average performance of user 4, and MOS outperforms round-robin by 22% while satisfying the long-term resource allocation requirement. Overall, the system performance obtained in MOS is only about 3% less than that of OSI, while both outperforming the round-robin by over 60% and satisfying the long-term resource

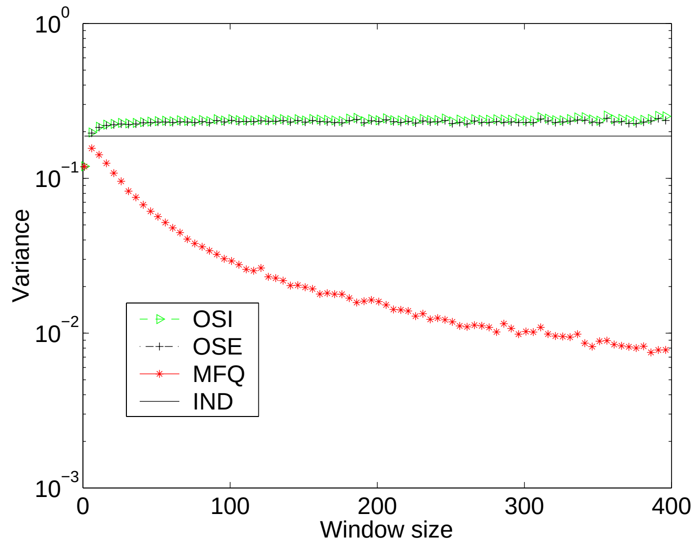

Fig. 3.9. Normalized variance of discrepancy as a metric of the short-term performance.

allocation requirement. Hence, MOS improves the short-term performance without dramatically decreasing the system throughput. In general, the larger the time-correlation, the worse the short-term performance, and the greater the improvement in the short-term performance, the larger the loss in system performance.

One closely related problem is to be able to handle users with explicit delay requirements, such as audio and video. It is a challenging problem to schedule users opportunistically while satisfying the delay requirements of certain users. One possible solution is to extend the current scheduling scheme in the following way. Each user has a due-time known by the base station. The due-time of a user is the deadline of the first packet in the user's queue. If a user has no delay requirement, its due-time is set to be ∞. Suppose different users have different due-times. (If two users have the same due-time, we randomly pick one and make its due-time the time-slot before its actual due-time.) We then adjust the transmission probability of a user of according to its due-time. The closer the due-time, the higher the probability of its transmission. When a user is at its due-time, we assign the time-slot to this user with probability 1. If we know the distribution functions of users' performance values, it is possible to determine how large the transmission probability should be (as a function of its and other users' due-times) numerically. Otherwise, these transmission probabilities might be determined experimentally.

3.7 Conclusion

In this chapter, we investigate an opportunistic transmission-scheduling with temporal fairness constraints as a case study. Given a temporal fairness requirement, the scheduling policy maximizes the average system performance. In our model, each user's performance value is a stochastic process, reflecting the time-varying performance that results from randomly-varying channel conditions. The users' performance-value processes can be arbitrarily correlated, both in time and across users. We establish the optimality of our opportunistic scheduling policy. We also provide a scheduling procedure that includes an on-line parameter-estimation algorithm to estimate parameter values used in the scheduling policy. Our scheduling algorithm has a low computational burden, which is important for on-line implementation. Via simulation, we illustrate the performance of our scheduling policy, showing significant performance gains over the round-robin policy. Our simulation results also show that our scheme works well for the case of nonstationary performance-value sequences, and is robust to estimation errors. Finally, we study a heuristic scheme to improve short-term performance.

The case study presented in this chapter is complete in the following sense: it includes an optimal scheduling policy, a parameter-estimation algorithm, implementation considerations and procedures, numerical results, consideration for estimation errors, and a heuristic scheme to improve short-term performance. Some of these issues, namely parameter estimation, implementation procedure, and considerations of estimation error and short-term performance, are common among other scheduling problems, which will be discussed in the next chapter.

4. A UNIFIED FRAMEWORK OF OPPORTUNISTIC SCHEDULING

In this chapter, we present a unified framework for opportunistically scheduling user transmissions to exploit the time-varying channel conditions in wireless communication systems. We consider a time-slotted system, such as TDMA or time-slotted CDMA. We use a stochastic model to capture the *time-varying* and *channel-condition-dependent* performance of each user. Specifically, $\{U_i^k\}$ is a stochastic process associated with user i, where U_i^k is the level of performance that would be experienced by user i if it is scheduled to transmit at time k. The value of U_i^k measures the "worth" or "utility" of time-slot k to the user i, and is in general a function of its channel condition. We assume that $\{U_i^k\}$ is stationary and ergodic, and U_i^k is *nonnegative and bounded*. We use the notation $\vec{U} = (U_1, \cdots, U_N)$, where U_i is a random variable representing the performance value of user i at a generic time-slot. (Note that the stationary assumption does not preclude correlations across users or across time.) Our goal in opportunistic scheduling is to maximize the average system performance under certain QoS requirements. The generic scheduling problem is as follows: given $\vec{U}$ (the vector of users' performance values), determine which user should be scheduled to transmit in the given time-slot, under the given QoS constraints.

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. We provide optimal scheduling solutions, and study the asymptotic behavior of our opportunistic scheduling schemes. We will also show how the case study presented in Chapter 3 and some previous work by other researchers directly fit into or is related to this framework (e.g., [47, 51, 64, 41]).

The chapter is organized as follows. In Section 4.1, we introduce a scheduling problem with temporal fairness requirements, which is a generalization of the problem discussed in Chapter 3, and provide an optimal solution. Then, we present a utilitarian fairness scheduling problem and its optimal solution in Section 4.2. Finally, a scheduling problem with minimum-performance guarantees and its solution are presented in Section 4.3. We compare different scheduling schemes in Section 4.4, and discuss an asymptotic result on opportunistic scheduling schemes in Section 4.5. We also show that how our results can be generalized to nonstationary policies and under more general conditions in Section 4.6. In Section 4.7, we briefly discuss the parameter estimation. In Section 4.8, we provide simulation results to illustrate the performance of the studied scheduling schemes, and present the conclusion in Section 4.9.

4.1 Temporal Fairness Scheduling Scheme

It is important to note that fairness criteria are central to scheduling problems in wireless systems. Without a good fairness criterion, the system performance can be trivially optimized by, for example, letting a user with the highest performance value to transmit. This may prevent "poor" users (in terms of either channel conditions or money) from accessing the network resource, and thus compromises the desirable feature of wireless networks: to provide "anytime," "anywhere" accessibility. In this chapter, we study two fairness criteria—temporal and utilitarian. In this section, we focus on the scheduling problem with temporal fairness requirements; we study utilitarian fairness in Section 4.2.

4.1.1 Problem Formulation

Recall that in Chapter 3, we study a scheduling problem with temporal fairness constraints; i.e., each user is assigned r_i portion of time-slots and $\sum_{i=1}^{N} r_i = 1$. We can further generalize it to give the system more freedom. To elaborate, let r_i denote the *minimum* time-fraction that should be assigned to user i, where $r_i \geq 0$, $\sum_{i=1}^{N} r_i \leq 1$. The new problem is more general, and allows more flexibility in resource sharing in that it allows for a minimum amount of fairness in the system. Note that $\epsilon := \sum_i r_i \leq 1$ is a tuning parameter such that the smaller the value of ϵ, the less restrictive the fairness constraint, and the greater the opportunity to improve the system performance. One extreme is $\epsilon = 1$, which is the case studied in Chapter 3. Another extreme is $\epsilon = 0$, where the system has the utmost freedom to maximize the system performance. Our goal is to develop a scheduling policy Q that exploits the time-varying channel conditions to maximize the total expected system performance while satisfying the (general) temporal fairness constraint. The problem can be stated formally as follows:

$$\begin{array}{lll} \underset{Q \in \Theta}{\text{maximize}} & E\left(U_{Q(\vec{U})}\right) & \\ \text{subject to} & P\{Q(\vec{U}) = i\} \geq r_i, & i = 1, 2, \cdots, N, \end{array} \tag{4.1}$$

where Θ is the set of all scheduling policies.

4.1.2 An Optimal Policy

We define a policy Q^* as follows:

$$Q^*(\vec{U}) = \underset{i}{\operatorname{argmax}}(U_i + v_i^*), \tag{4.2}$$

where the v_i^*s are chosen such that:

1. $\min_i(v_i^*) = 0$
2. $P\{Q^*(\vec{U}) = i\} \geq r_i$ for all i
3. For all i, if $P\{Q^*(\vec{U}) = i\} > r_i$, then $v_i^* = 0$.

Proposition 3 *The policy Q^* is a solution to the problem defined in (4.1), i.e., it maximizes the average system performance under the temporal fairness constraint.*

Proof: Let Q be a policy satisfying $P\{Q(\vec{U}) = i\} \geq r_i$ for all i. Also recall that $v_i^* \geq 0$. Hence, we have

$$\begin{aligned}
E\left(U_{Q(\vec{U})}\right) &\leq E\left(U_{Q(\vec{U})}\right) + \sum_{i=1}^{N} v_i^*(P\{Q(\vec{U}) = i\} - r_i) \\
&= E\left(\sum_{i=1}^{N} U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right) + \sum_{i=1}^{N} v_i^*(P\{Q(\vec{U}) = i\} - r_i) \\
&= E\left(\sum_{i=1}^{N} (U_i + v_i^*)\, \mathbf{1}_{\{Q(\vec{U})=i\}}\right) - \sum_{i=1}^{N} v_i^* r_i.
\end{aligned}$$

(The notation $\mathbf{1}_{\{Q(\vec{U})=i\}}$ represents the indicator function of event $\{Q(\vec{U}) = i\}$.) By the definition of Q^*, we have $\sum_{i=1}^{N}(U_i + v_i^*)\mathbf{1}_{\{Q(\vec{U})=i\}} \leq \sum_{i=1}^{N}(U_i + v_i^*)\mathbf{1}_{\{Q^*(\vec{U})=i\}}$. Thus,

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

Hence,

$$\begin{aligned}
E\left(U_{Q(\vec{U})}\right) &\leq E\left(\sum_{i=1}^{N} (U_i + v_i^*)\, \mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) - \sum_{i=1}^{N} v_i^* r_i \\
&= E\left(U_{Q^*(\vec{U})}\right) + \sum_{i=1}^{N} v_i^*(P\{Q^*(\vec{U}) = i\} - r_i) \\
&= E\left(U_{Q^*(\vec{U})}\right),
\end{aligned}$$

which completes the proof. □

We can think of the parameter $\vec{v}^*$ in (4.2) as an "offset" used to satisfy the fairness requirement. Consider the case where we want to maximize the overall performance without any QoS requirements. It is straightforward to show that we should always choose the "best" user (i.e., the user with the maximum performance value) to transmit in a generic time-slot. In other words, $Q(\vec{U}) = \operatorname{argmax}_i U_i$. However, such a scheme may be unfair to certain users. Hence, to satisfy the fairness requirement, the scheduling policy schedules the "relatively-best" user to transmit. User i is "relatively-best" if $U_i + v_i^* \geq U_j + v_j^*$ for all j. If $v_i^* > 0$, then user i is an "unfortunate" user, i.e., the channel condition it experiences is relatively poor. Hence, it has to take advantage of some other users (e.g., users with $v_j^* = 0$) to satisfy its fairness requirement. Thus, to maximize the overall system performance, we can only give the "unfortunate" users the amount of resource equivalent to their minimum requirements. When $P\{Q^*(\vec{U}) = j\} > r_j$ for user j, the user gets more than its minimum requirement—this user cannot take advantage of other users, i.e., $v_j^* = 0$.

The values of v_i^* are determined by the distribution of $\vec{U}$ and the values of r_i. In practice, the distribution of $\vec{U}$ is unknown, and hence we need to estimate the parameters v_i^*. Similarly, in the opportunistic scheduling schemes discussed in Sections 4.2 and 4.3, there are also parameters that need to be estimated. In Section 4.7, we explain how to use stochastic approximation algorithms to efficiently estimate the values of these parameters. The algorithm is also based on stochastic approximation as studied in Chapter 3.

The policy Q^* maximizes the average system performance even if the users' performance values are arbitrarily correlated, both in time and across users. The following proposition establishes, under a more restrictive assumption, that each user can be guaranteed a minimum performance level given the distribution of $\vec{U}$.

Proposition 4 *If the performance values U_i, $i = 1, \cdots, N$, are independent, then for all i,*

$$E\left(U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) \geq P\{Q^*(\vec{U}) = i\}E(U_i) \geq r_i E(U_i).$$

An alternative statement of the above result is

$$E(U_i | Q^*(\vec{U}) = i) \geq E(U_i).$$

The proposition is a trivial generalization of Prop. 2. For a proof of the proposition, see Appendix A.2. Note that $E(U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}})$ is the average performance value of user i using our opportunistic scheduling policy, and $P\{Q^*(\vec{U}) = i\}E(U_i)$ is the average performance of user i when using a non-opportunistic scheduling scheme where $P\{Q^*(\vec{U}) = i\}$ portion of the resource (i.e., time) is assigned to user i. (A non-opportunistic scheduling policy is one that does not use information on channel conditions to decide which user to transmit.)

The above proposition guarantees, assuming the users' performance values are independent, that the average performance of *each user* in our opportunistic scheduling scheme will be no worse than that of any non-opportunistic scheduling scheme that allocates the same share of the resource to the user. Furthermore, each user gets a guaranteed minimum performance of $r_i E(U_i)$ because $P\{Q(\vec{U}) = i\} \geq r_i$. This result is intuitively appealing. When a user is experiencing good channel conditions, it has a higher chance to have a maximum value of $U_i + v_i^*$ among all users, and thus be chosen to transmit. When a user is experiencing poor channel conditions, it has less opportunity to be the relatively-best user and thus to be scheduled. Hence, a user tends to transmit more often under favorable conditions, resulting in performance improvement for each user. In this sense, the opportunistic scheduling policy gives all the users the chance to improve their expected performance. Of course, different users may experience different levels of improvement. In general, the larger the variability of a user's performance value, the greater the improvement.

4.2 Utilitarian Fairness Scheduling Scheme

In the last section, we studied the opportunistic scheduling problem with temporal fairness requirements. In wireline networks, when a certain amount of resource is assigned to a user, it is equivalent to granting the user a certain amount of throughput/performance value. However, the situation is different in wireless networks, where the amount of resource and the performance value are not directly related (though correlated). Hence, in this section we describe an alternate scheduling problem that would ensure that all users get at least a certain portion of the system performance.

4.2.1 Problem Formulation

Recall that $E(U_i\mathbf{1}_{\{Q(\vec{U})=i\}})$ is the average performance value of user i using policy Q, and $E(U_{Q(\vec{U})}) = \sum_{i=1}^{N} E(U_i\mathbf{1}_{\{Q(\vec{U})=i\}})$ is the overall performance of the system using policy Q. Let a_i be the minimum fraction of the overall average performance required by user i, where $a_i \geq 0$ for all i, and $\sum_i a_i \leq 1$. Then the optimal opportunistic scheduling problem based on utilitarian fairness can be written as:

$$\begin{array}{lll} \underset{Q\in\Theta}{\text{maximize}} & E\left(U_{Q(\vec{U})}\right) & \\ \text{subject to} & E\left(U_i\mathbf{1}_{\{Q(\vec{U})=i\}}\right) \geq a_i E\left(U_{Q(\vec{U})}\right), & i = 1, 2, \cdots, N, \end{array} \tag{4.3}$$

where Θ is the set of all policies. Note that a_is are predetermined fairness parameters, and $\epsilon' = \sum_i a_i$ is a tuning parameter (similar to ϵ in Section 4.1)—the smaller its value, the larger the opportunity to improve system performance.

The authors of [51] consider a special case of the above opportunistic scheduling problem. Specifically, they consider maximizing the minimum weighted performance of users (i.e., $E(U_i\mathbf{1}_{\{Q(\vec{U})=i\}})/\beta_i$, where the β_is are predetermined weights). This is then equivalent to setting $a_i = \beta_i / \sum_j \beta_j$ in the utilitarian fairness problem defined by (4.3). In this case, because $\sum_i a_i = 1$, all the inequality constraints are active for any feasible solution, i.e., the inequality constraints are satisfied with equality.

This problem setting requires fairness in terms of performance values, which, to some extent, parallels the concept of weighted fair queueing used in wireline networks. The difference is that the overall capacity here is not fixed; it depends on channel conditions, the values of a_i, and the scheduling policy.

4.2.2 An Optimal Policy

We define a policy Q^* as:

$$Q^*(\vec{U}) = \underset{i}{\operatorname{argmax}} \left((\kappa + \nu_i^*) \, U_i \right), \tag{4.4}$$

where $\kappa = 1 - \sum_{i=1}^{N} a_i \nu_i^*$, and the ν_i^*s are chosen so that:

1. $\min_i(\nu_i^*) = 0$

2. $E\left(U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) \geq a_i E\left(U_{Q^*(\vec{U})}\right)$ for all i

3. For all i, if $E\left(U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) > a_i E\left(U_{Q^*(\vec{U})}\right)$, then $\nu_i^* = 0$.

Proposition 5 *The policy Q^* in (4.4) is a solution to the problem defined in (4.3), i.e., it maximizes the average system performance under the utilitarian fairness constraint.*

For a proof of the above proposition, see Appendix A.3. Similar to $\vec{v}^*$ in the last section, the parameter $\vec{\nu}^*$ in (4.4) can be considered a "scaling" used to satisfy the utilitarian fairness constraint. The optimal scheduling policy always chooses the relatively-best user to transmit. In this case, the user i is relatively-best if $(\kappa + \nu_i^*) U_i = \max_j (\kappa + \nu_j^*) U_j$, where κ is a constant for all users. As before, if $\nu_i^* > 0$, then user i is an "unfortunate" user, and its average performance value equals the minimum requirement, i.e., $E(U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}}) = a_i E(U_{Q^*(\vec{U})})$.

Utilitarian scheduling schemes have certain features. First, any policy Q that satisfies the fairness constraint defined in (4.3) has the property that

$$\frac{a_i}{1 - \epsilon' + a_j} \leq \frac{E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right)}{E\left(U_j \mathbf{1}_{\{Q(\vec{U})=j\}}\right)} \leq \frac{1 - \epsilon' + a_i}{a_j},$$

for $i, j = 1, 2, \cdots, N$. In other words, the utilitarian fairness constraint controls the maximum discrepancy of performance values among users.

Second, the constraint given in (4.3) ensures that a user is given at least a certain share of the total performance, and is hence more suitable in some situations than the temporal fairness constraint given by (4.1). However, there is also a significant disadvantage of a utilitarian scheduling scheme: a user experiencing poor channel conditions could have a detrimental impact on the overall system performance. By observing the constraint in (4.3), we have

$$E\left(U_{Q(\vec{U})}\right) \leq \frac{E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right)}{a_i} \leq \frac{E(U_i)}{a_i}.$$

Hence, if a user is experiencing very poor channel conditions [a very small value of $E(U_i)$] and has a large value of a_i, then it could compromise the overall system performance significantly because a substantial portion of the total time-slots may have to be allocated to this user in order to meet its fairness requirement. To alleviate this potential problem, one can devise an adaptive thresholding strategy. To elaborate, if

$$\frac{E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right)}{P\{Q(\vec{U})=i\}E\left(U_{Q(\vec{U})}\right)} \leq \beta,$$

where β is a predetermined threshold, then we decrease the values of a_i because user i cannot utilize the scarce spectrum efficiently.

We have studied two different fairness criteria—temporal and utilitarian fairness. In the temporal fairness scheme, users "interact" with each other through resource sharing. Users' behavior are relatively isolated; i.e., given r_i (the minimum time-fraction assigned to user i), the achieved performance of a user depends heavily on its own performance values. A very poor user can at most waste r_i portion of the system resource. This is different in utilitarian fairness schemes, where the achieved performance values of users are heavily correlated because each user shares a certain percentage of the overall performance.

4.3 Minimum-Performance Guarantee Scheduling Scheme

Thus far, we have discussed two optimal scheduling schemes that provide users with different fairness guarantees. However, while they satisfy a relative measure of performance (i.e., fairness), they do not consider any absolute measures. This motivates the study of a new category of scheduling problems where QoS is defined as the minimum-performance guarantee. To elaborate, the objective is to maximize the average system performance subject to meeting each user's minimum-performance requirement.

4.3.1 Problem Formulation

Suppose that each user has a minimum-performance requirement C_i, and the vector $\vec{C} = \{C_1, C_2, \cdots, C_N\}$ is the *requirement vector*. The problem to maximize the system performance while satisfying each user's minimum requirement is stated as:

$$\begin{aligned} \underset{Q}{\text{maximize}} \quad & E\left(U_{Q(\vec{U})}\right) \\ \text{subject to} \quad & E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right) \geq C_i, \qquad i = 1, 2, \cdots, N. \end{aligned} \tag{4.5}$$

Consideration of this problem raises two questions: (i) Is $\vec{C}$ a feasible requirement vector, i.e., does there exist a policy Q such that $E(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}) \geq C_i$ for all i? (ii) If $\vec{C}$ is a feasible requirement vector, which policy maximizes the overall performance under the given QoS requirement?

Unlike in our previous two problems, the QoS constraint in this problem is defined as a minimum-performance requirement instead of a fairness requirement. Hence, the formulation here offers users a more "direct" service guarantee. For example, if the performance measure is defined as the data-rate, then each user is guaranteed a minimum data-rate, which may be more important to a user than knowing that a minimum amount of resource will be assigned to it. While appealing to users, providing minimum-performance guarantees can be quite difficult in practice because of the feasibility issue—can the system satisfy the performance requirements for all users? Note that feasibility is not a concern in the fairness-based constraints. In the temporal-fairness scheduling problem defined in (4.1), as long as $\sum_i r_i \leq 1$, the system is feasible. In the utilitarian fairness scheduling problem defined in (4.3), $\sum_i a_i \leq 1$ is the feasibility constraint. Both of them are easy to verify. However, there is no easy way, in general, to determine whether a given $\vec{C}$ is feasible or not. We discuss this issue in more detail in Section 4.3.3.

There are, however, some natural settings where feasibility is not a problem. For example, let $C_i = \rho_i E(U_i)$, where $\rho_i \geq 0$ for all i and $\sum \rho_i \leq 1$. Such a setting can be satisfied by a non-opportunistic scheduling policy in which user i is chosen to transmit in a given time-slot with probability ρ_i. It is hence feasible for opportunistic scheduling policies.

4.3.2 An Optimal Policy

We present an optimal scheduling policy assuming feasibility in this section, and then discuss the feasibility problem in Section 4.3.3. Suppose $\vec{C} = \{C_1, C_2, \cdots, C_N\}$ is feasible. We define a policy Q^* by:

$$Q^*(\vec{U}) = \underset{i}{\operatorname{argmax}}(\alpha_i^* U_i), \tag{4.6}$$

where the α_i^*s are chosen so that:

1. $\min_i(\alpha_i^*) = 1$
2. $E\left(U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) \geq C_i$ for all i
3. For all i, if $E\left(U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) > C_i$, then $\alpha_i^* = 1$.

Proposition 6 *The policy Q^* defined in (4.6) is a solution to the problem defined in (4.5), i.e., it maximizes the average system performance under the minimum-performance constraint.*

For a proof of the above proposition, see Appendix A.4. The parameter $\vec{\alpha}$ "scales" the performance values of users, and the scheduling policy schedules the relatively-best user, where user i is relatively-best if $\alpha_i^* U_i = \max_j \alpha_j^* U_j$. If the scaling factor for a user is larger than 1, then the user is an "unfortunate" user, and it is granted only an average performance value that equals its minimum performance requirement.

Our opportunistic scheduling policy dominates non-opportunistic policies in the following sense. Consider a non-opportunistic scheduling policy in which user i shares a portion ρ_i of the resource (time-slots), where $\sum_i \rho_i = 1$, and user i gets an average performance value $\rho_i E(U_i)$. Let $C_i = \rho_i E(U_i)$ for all i. Then $\vec{C}$ is feasible, and the opportunistic scheduling policy always provides "no-worse" performance values for each user relative to that of the non-opportunistic scheduling policy, assuming that the signaling cost is negligible.

In [46, 47], the authors study scheduling algorithms where both delay and channel conditions are taken into account. Roughly speaking, the algorithm is: argmax $b_i W_i t_i$, where W_i is the head-of-the-line packet delay for queue i, t_i is the channel capacity, and b_i is some constant. Furthermore, the authors of [47] state the following result (assuming there is a finite set of channel states): to maximize the system throughput with minimum-throughput requirements, there exists some constant c_i such that one should choose a user with the maximum value of $c_i t_i$. In these papers, however, there is no discussion on how to obtain the values of c_i, how to break ties, or how feasibility plays a role. These issues are addressed in this section.

It turns out that a scheduling policy of the form of $Q^* = \operatorname{argmax} \alpha_i U_i$ (as in (4.6)) is optimal for other types of scheduling problems as well. For example, the optimal utilitarian policy defined in (4.4) is in this form (i.e., $\alpha_i = \kappa + \nu_i^*$). Moreover, consider a class of scheduling problems without explicit constraints. Let u_i be the data rate of user i at a generic time-slot (an example of the performance value), and let $\vec{u} = (u_1, \cdots, u_N)$. Then, the average data rate of user i using policy Q, $t_i(Q)$, is given by $t_i(Q) = E\left(u_i \mathbf{1}_{\{Q(\vec{u})=i\}}\right)$. Suppose $f_i(x)$ is a monotonically increasing function of x, representing the utility of user i given data-rate x. The problem is to find a policy that maximizes the overall system utility:

$$\underset{Q}{\text{maximize}} \sum_i f_i(t_i(Q)). \tag{4.7}$$

Then following Prop. 6, an optimal policy Q^* with respect to (4.7) is $Q^*(\vec{u}) = \operatorname{argmax}_i(\alpha_i u_i)$, where $\vec{\alpha}$ is a set of parameters that depend on the functions f_i and distribution functions of U_i. Different scheduling policies will result from different choices of the function f_i. Some examples are provided next.

The first example is when $f_i(x) = \log(x)$ and $\alpha_i = 1/E(u_i \mathbf{1}_{\{Q(\vec{u})=i\}})$, which result in the well-known proportional-fairness scheduling algorithm described in [41]. The second example is to maximize the overall throughput where $f_i(x) = x$. It is obvious that $\vec{\alpha} = [1, 1, \cdots, 1]$ is the solution to this problem, i.e., we always choose the user with the highest data-rate to transmit. The last example is our scheduling problem with minimum-throughput guarantees. In this case,

$$f_i(x) = \begin{cases} -\infty & \text{if } x < C_i \\ x & \text{if } x \geq C_i \end{cases},$$

and an appropriate $\vec{\alpha}$ is given in (4.6). In summary, the solution to a very large group of opportunistic scheduling problems is in the form of $\operatorname{argmax} \alpha_i U_i$.

4.3.3 Feasibility

The ability to provide a specific performance guarantee is an advantage of our scheme. However, to satisfy such a constraint introduces the question of feasibility. In the following, we discuss the feasibility problem and how to determine the feasibility region of our scheduling policy. The feasibility region of a policy Q is the set of requirement vectors that are feasible under the policy Q.

Proposition 7 *The feasible region of our opportunistic scheduling policy is* convex *and contains the feasibility regions of all policies.*

A proof is included in Appendix A.6. The feasibility region of our scheduling policy is determined by the distribution of $\vec{U}$. In general, there is no closed-form expression for the feasibility region even if the distribution function of $\vec{U}$ is known. The distribution of $\vec{U}$ depends on the user's channel condition, its mobility, and the form of the performance value function. It is practically impossible to know the distribution of $\vec{U}$ *a priori* in the system. Hence, we estimate the feasibility region using sample paths; i.e., the sequence of $\{U_i^k\}$. Convexity is an important feature in determining whether a requirement vector is feasible.

In general, the feasibility region is in an N-dimensional space, where N is the number of users in the system. The vertex on the ith axis is $[0, 0, \cdots, E(U_i), \cdots, 0]$, which corresponds to assigning all the resource to user i. In the extreme case that all the resource is assigned to a single user, there is no performance difference between opportunistic and non-opportunistic scheduling policies. Hence, they share the same vertices on the axes. These N vertices span an $(N-1)$-dimensional hyperplane. Any *non-negative* vector below this hyper-plane is a feasible requirement vector for a non-opportunistic scheduling policy.

The feasibility region of the opportunistic scheduling policy contains the feasibility region of any non-opportunistic scheduling policy, while they share the same vertices on the axes. We next determine some other vertices in the feasibility region of the opportunistic scheduling policy. Let

$$Q_{\vec{\alpha}}(\vec{U}) = \underset{i}{\operatorname{argmax}}(\alpha_i U_i), \tag{4.8}$$

where ties are broken randomly. Given a value of $\vec{\alpha}$, using policy $Q_{\vec{\alpha}}(\vec{U})$ results in an average performance vector, where its ith component is the average performance value of user i (i.e., $E(U_i \mathbf{1}_{\{Q_{\vec{\alpha}}(\vec{U})=i\}})$). In other words, by choosing a value for the vector $\vec{\alpha}$, we obtain an average performance-value vector that determines one point on the boundary of the feasibility region. By varying the values of $\vec{\alpha}$, we can draw the boundary of the feasibility region. For example, if we set $\vec{\alpha} = [1, 1, \cdots, 1]$ in (4.8), then we get the average performance-value vector representing the maximum performance the system can obtain. By using different values of $\vec{\alpha}$ in (4.8), we get different performance-value vectors, resulting in different points in the N-dimensional space. These points, along with the N vertices in the N axes, span an N-dimensional surface. Because the feasibility region is *convex*, any *non-negative* vector under this surface is feasible. If we choose more values of $\vec{\alpha}$, we get more points on the boundary of the feasibility region, and thus we get a closer approximation to the actual feasibility set.

Figure 4.1 shows the feasibility region for two users. The performance values of user 1 and user 2 are independent and exponentially distributed with mean values 4 and 5, respectively. The two vertices on the two axes correspond to the two extreme cases that all the resource is assigned to one user. The area between the straight line (a

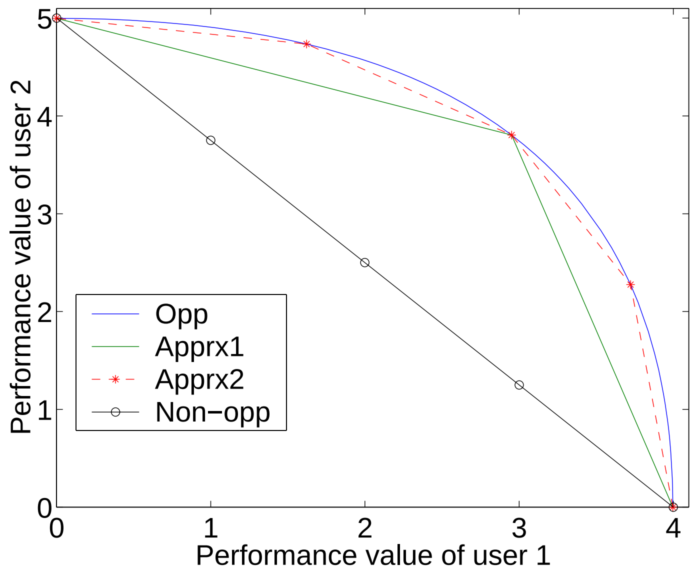

Fig. 4.1. The feasible region of two users.

1-dimensional "plane") and the two axes is the feasibility region of a non-opportunistic scheduling policy. The uppermost solid curve indicates the boundary of the feasibility region of the opportunistic scheduling policy. The area between the uppermost curve and the two axes is the feasibility region of our opportunistic scheduling policy. The other two curves in between are approximations of the boundary of the feasibility region. The solid line is the case where we have one vertex ($\vec{\alpha} = [1, 1]$)) besides the two vertices on the axes. Note that the vertex corresponding to $\vec{\alpha} = [1, 1]$ is the maximum-system-performance case. By adding two more points ($\vec{\alpha} = [1, 3]$ and $\vec{\alpha} = [3, 1]$), we get a better approximation (dashed curve with stars), which gives a fairly close approximation of the feasibility region. Hence, it is fair to say that we can obtain reasonable estimates the feasibility region via measurement data.

Next, we discuss how to determine the feasibility region when the set of (active) users changes in the cell. Suppose there are currently N users in the system, and we have some information about these N users. We can estimate the feasibility region for the N users using different values of $\vec{\alpha}$. Note that we can draw these vertices simultaneously via measurements of $\vec{U}$. We consider two situations: First, if a user leaves the system, we simply collapse the feasibility set from N dimensions to $N-1$ dimensions by removing the axis of the leaving user. Second, suppose a new user joins the system, and we do not have any information on the user except its average performance value $E(U_{N+1})$. We can connect all the points on the surface of the feasibility set for the original N users with the new vertex on the $(N+1)$th axis, and construct a surface of $N+1$ dimensions. Due to the convexity, any non-negative $(N+1)$-dimensional vector under the new surface is feasible.

Compared with non-opportunistic schedulers, opportunistic schedulers *enlarge* the feasible region/capacity. Hence, we can achieve the following goals:

1. Accommodate more users under the same admission control policy. For example, consider a simple admission control algorithm—users are admitted in the system as long as the requirement vector is feasible; i.e., the admission region is the same as the feasibility region. As shown in Figure 4.1, the uppermost curve indicates the boundary of the feasibility region of the opportunistic scheduling policy, whereas the straight (lowest) line is that of a non-opportunistic scheduling policy. The area between these two curves is the area where the opportunistic scheduling policy can accommodate two users while the non-opportunistic scheduling policy can only accommodate one user.

2. Improve users' service quality, in terms of higher performance and/or lower degradation probability, when the same users are admitted in a non-opportunistic scheduling scheme. Degradation means that the system cannot satisfy the performance "guarantee" given to a user at admission due to system load changes (say, handoff), variation in channel conditions, and mobility. If we choose feasible C_i and let $C_i \geq r_i E(U_i)$, then our opportunistic scheduling policy results in a "no-loss" situation for each user over non-opportunistic scheduling policies, which guarantees better service quality for users.

4.4 Discussion

We have presented a framework for opportunistic scheduling and studied three classes of scheduling problems under the framework. These scheduling problems share a common goal: to improve the spectrum efficiency while maintaining certain levels of QoS for each user using opportunistic scheduling algorithms. The solutions to these scheduling problems also have certain similarities—all the schemes choose the "relatively-best" user to transmit. Although "relatively-best" has a different meaning for each scheduling policy, the basic idea is to use an off-set or a scaling to satisfy the QoS requirements for users. If a user is "unfortunate", i.e., it has to take advantage of other users to satisfy its QoS requirement (in terms of fairness or performance value), then the user does not get more than its minimum requirement. This is done to maximize the system performance under the given constraints. In general, the larger the number of users sharing the same channel, or the larger the variance of $\vec{U}$, the larger the "opportunistic" scheduling gain compared with non-opportunistic scheduling policies. Furthermore, the more restrictive the QoS constraint, the less the flexibility for opportunistic scheduling decisions, and the lower the system performance gain.

In this chapter, we consider fairness from two different aspects: temporal fairness and utilitarian fairness. Max-min fairness [65] can be considered as a special case of the fairness requirement presented here. The intuitive notion of max-min fairness is that any user is entitled to as much performance/resource as any other user. Max-min fairness can be applied in two different ways. First, if we apply max-min fairness to the utilitarian fairness scheduling scheme, then the system should maximize the minimum performance of the users. This is equivalent to setting $a_i = 1/N$ for all i in (4.3) (N is the number of users sharing the same channel). Specifically, each user obtains the same performance value, and the system maximizes it. Second, if max-min fairness is applied to the temporal fairness scheduling scheme, then we have $r_i = 1/N$ for all i in (4.1), i.e., each user is granted to the same amount of resource.

As mentioned earlier, feasibility is not a concern in the scheduling problems with fairness requirements. The feasibility issue only arises in the problem with minimum-performance requirements.

The scheduling schemes with temporal fairness requirements or minimum-performance requirements can guarantee that the performance of each user is at least as good as that of the corresponding user in any non-opportunistic scheduling scheme under some assumptions. This desirable property cannot be guaranteed for scheduling schemes with utilitarian fairness requirements.

Different schemes may be suitable for different scenarios. For example, if the service provider wants to build a simple wireless network with pricing, the temporal fairness scheduling scheme is a reasonable choice. The temporal fairness scheduling scheme is simple and flexible without feasibility concerns. The amount of resource consumed by a user implies the minimum performance the user gets (with technical assumptions, see Prop. 4). The resource consumed by a user can be connected directly with the price the user should pay. On the other hand, the minimum-performance guarantee scheme provides users a direct performance assurance, but involves the additional complication of feasibility. If the service provider wants to build a network that provides data-rate guarantees, then this scheme is an appropriate choice. However, in practice, the feasibility issue may be difficult to handle especially in a wireless setting, and providing service performance guarantees is challenging in both wireless and wireline networks.

It should be noted that our framework for opportunistic scheduling can also cover cases where there are different constraints from different users. For example, some users may have resource requirements while other users can have a minimum-performance requirements. In such scenarios, similar optimal solutions can be provided under our framework using similar optimization techniques.

Last, we should note that the problem formulations, the objectives, and the constraints are expressed in terms of "expectation," which is a "long-term" performance measure. There is no guarantee of "short-term" performance. In Chapter 3 and in [64], we discuss an extension to improve short-term performance, and a similar process can be applied to other scheduling schemes. The basic idea is to increase a user's probability of transmission when it is behind in its share. In general, when users' performance values have strong correlation across time, the short-term performance is poor. The stricter the short-term performance requirement, the lower the opportunity to exploit time-varying channel conditions, and the less the performance improvements. We also refer interested readers to [46, 47, 49, 48] where queueing delays are considered.

4.5 Asymptotic Performance Bound

In this section, we study the asymptotic behavior of the opportunistic scheduling schemes. Consider first the greedy scheduling policy $Q(\vec{U}) = \operatorname{argmax}_i U_i$. This policy always chooses the best user to transmit, and thus achieves the highest system performance among *all* scheduling policies. Let

$$Z_n = U_{Q(\vec{U})} = \max_{i=1,\cdots,n} U_i, \tag{4.9}$$

where n is the number of users. Then $E(Z_n)$ is the average system performance of policy Q, the maximum among scheduling policies. Note that $E(Z_n)$ is a *tight* upperbound on the performance of all the opportunistic scheduling schemes studied in the dissertation (tight in the sense that there are scheduling schemes that come arbitrarily close to the bound, and there exist degenerate cases where the optimal scheme is the greedy scheme and hence achieves the bound). For example, if $\alpha_i = 0$ for all i in the temporal fairness scheduling scheme, then $Q(\vec{U}) = \operatorname{argmax}_i U_i$ is an optimal solution following Prop. 3. Similarly, when $a_i = 0$ and $C_i = 0$ for all i, the greedy algorithm is an optimal solution in the utilitarian fairness and minimum-performance guarantee schemes, respectively. From (4.9), it is obvious that $Z_{n+1} \geq Z_n$, i.e, the average system performance increases as the number of users competing for the same channel increases. But how fast can $E(Z_n)$ increase? The following result gives us an upperbound.

Proposition 8 *Suppose that $E(|U_i|) \leq C < \infty$ for all i. Then*

$$E(Z_n) = O(n), \tag{4.10}$$

and $O(n)$ is a "tight" bound in the following sense: for any $\epsilon > 0$, there exists a sequence of random variables $\{U_i\}$ such that

$$\lim_{n\to\infty} \frac{E(Z_n)}{n^{1-\epsilon}} = \infty.$$

A proof of this proposition is provided in Appendix A.7. By studying the asymptotic behavior of $E(Z_n)$, we obtain insights on the potential (limit) of opportunistic scheduling algorithms. As presented in the proposition, $E(Z_n)$ grows at most as fast as $O(n)$, and $O(n)$ is a "tight bound" in the sense that there exists sequences of random variables that can reach this bound arbitrarily closely. Of course, the behavior of $E(Z_n)$ is different when the U_is have different distribution functions. We next illustrate the maximum achievable performance in opportunistic scheduling schemes for various U_i. For simplicity, in each case, we assume that U_i are i.i.d. (independent and identically distributed) random variables. In this case, any non-opportunistic scheduling scheme obtains a system performance value of $E(U)$, which is not affected by the number of users in the system, where U is a random variable with the same distribution as U_i. We define a ratio $G_n = E(Z_n)/E(U)$ to illustrate the performance gain of an opportunistic scheduling scheme over that of non-opportunistic ones.

- If U_i is uniformly distributed over an interval $[a, b]$, then

  $$\lim_{n\to\infty} E(Z_n) = b,$$

  and $\lim_{n\to\infty} G_n = 2$.

- If U_i is exponentially distributed with mean θ, then

  $$E(Z_n) = \theta(1 + \sum_{i=1}^{n-1} \frac{1}{i}).$$

  Hence, $\lim_{n\to\infty} G_n / \log n = 1$.

- Let

$$F(u) = \begin{cases} 1 - \frac{1}{u^\alpha} & u \geq 1 \\ 0 & u < 1. \end{cases}$$

and $\alpha > 1$. Then $E(U) = \frac{1}{\alpha-1} < \infty$, and

$$\lim_{n\to\infty} \frac{E(Z_n)}{n^{1/\alpha}} = E_0,$$

where $E_0 = \int_0^\infty 1 - \exp(-x^{-\alpha})\, \mathrm{d}x$. We have $\lim_{n\to\infty} G_n/n^{1/\alpha} = E_0(\alpha - 1)$, and $E(Z_n)$ gets close to the bound $O(n)$ as α decreases ($\alpha > 1$).

4.6 Generalization

Note that in the previous problem formulations we only consider stationary policies. (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 our result to more general cases. Similar extensions hold for other scheduling problems.

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

$$\begin{aligned} F_i^K(Q) &= \frac{1}{K}\sum_{k=1}^{K} U_i^k \mathbf{1}_{\{\mathcal{Q}^k=i\}}, \quad i = 1, 2, \cdots, N \\ R_i^K(Q) &= \frac{1}{K}\sum_{k=1}^{K} \mathbf{1}_{\{\mathcal{Q}^k=i\}}, \quad i = 1, 2, \cdots, N, \end{aligned}$$

where $\mathcal{Q}^k = Q(\{\vec{U}^t, t = 1, 2, \cdots\}, k)$. 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

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

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

We restate our policy Q^* as follows:

$$Q^*(\vec{U}^k) = \underset{i}{\text{argmax}}(U_i^k + v_i^*), \tag{4.11}$$

where the v_i^*s are chosen such that:

1. $\min_i(v_i^*) = 0$
2. $\liminf_{K\to\infty} R_i^K(Q^*) \geq r_i$ for all i
3. For all i, if $\liminf_{K\to\infty} R_i^k(Q^*) > r_i$, then $v_i^* = 0$.

Let Θ be the set of all scheduling policies. The temporal scheduling problem is formulated as:

$$\begin{array}{ll} \underset{Q\in\Theta}{\text{maximize}} & F(Q) \\ \text{subject to} & \liminf\limits_{K\to\infty} R_i^K(Q) \geq r_i, \qquad i = 1, 2, \cdots, N. \end{array} \tag{4.12}$$

Proposition 9 *If* $\lim_{K\to\infty} R_i^K(Q^*)$ *exists for all* i *for the* Q^* *defined in (4.11), then the policy* Q^* *is a solution to the problem defined in (4.12).*

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

- Suppose that $\{\vec{U}^k, k = 1, 2, \cdots\}$ is stationary and ergodic. Because Q^* is a stationary policy, $R_i^K(Q^*)$ and $F_i^K(Q^*)$ converge to a constant almost surely. Thus, Q^* is a solution to the problem defined in (4.12). Furthermore, we have

$$\liminf_{K\to\infty} F^K(Q^*) = \limsup_{K\to\infty} F^K(Q^*). \tag{4.13}$$

  This equation is critical. It states the important fact that the asymptotic worst-case system performance of our policy Q^* ($\liminf_{K\to\infty} F^K(Q^*)$) is the same as its asymptotic best-case system performance ($\limsup_{K\to\infty} F^K(Q^*)$). 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 $\{\vec{U}^k, k = 1, 2, \cdots\}$ is stationary and ergodic. Then many policies have the nice property that $F_i^K(Q)$ and $R_i^K(Q)$ converge (to a random variable almost surely). Examples of such policies are stationary policies and periodic policies[1]. However, there exist policies such that

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

  In this case, even if the policy Q is a solution to the problem defined in (4.12), 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 (4.12)) bounds the asymptotic best-case performance of others.

- The proposition holds without the assumption that $\{\vec{U}^k, k = 1, 2, \cdots\}$ is stationary and ergodic. However, in this case, we may not be able to estimate the parameters (v_i^*) used in Q^* in practice.

- In the dissertation, we often use round-robin policy as an example of non-opportunistic policies for comparison. To be specific, round-robin is a non-stationary non-opportunistic scheduling policy. If $\{\vec{U}^k, k = 1, 2, \cdots\}$ 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., $\liminf_{K\to\infty} R_i^K(Q) \geq r_i$ for all i, then for any $\epsilon > 0$, there exists L_1, such that for any $K > L_1$, we have

$$R_i^K(Q) > r_i - \frac{\epsilon}{2\sum_{i=1}^N v_i^*}, \quad i = 1, 2, \cdots, N. \tag{4.14}$$

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

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

$$v_i^*\left(\lim_{K\to\infty} R_i^K(Q^*) - r_i\right) = 0, \quad i = 1, 2, \cdots, N.$$

Hence, for the ϵ in (4.14), there exists $L > L_1$, such that for $K > L$, we have

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

Then for $K > L$, we have

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

By the definition of Q^*, we have $\sum_{i=1}^N (U_i^k + v_i^*)\mathbf{1}_{\{\mathcal{Q}^k=i\}} \leq \sum_{i=1}^N (U_i^k + v_i^*)\mathbf{1}_{\{Q^*(\vec{U}^k)=i\}}$. Thus,

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

Because ϵ is chosen arbitrarily, we have

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

which complete the proof. $\square$

4.7 Implementation

The opportunistic scheduling policies described in previous sections all involve some parameters that need to be estimated online. For example, the temporal fairness scheduling policy is given by

$$Q^*(\vec{U}) = \underset{i}{\operatorname{argmax}}(U_i + v_i^*),$$

where the v_i^*s are parameters determined by the distribution of $\vec{U}$ and values of the r_i. In practice, this distribution is *a priori* unknown, and hence we need to estimate the parameters v_i^*. Similar to the case study presented in Chapter 3, we use a stochastic-approximation-based algorithm to estimate the parameters. The difference is due to the inequality constraints. In the following, we use the (general) temporal fairness scheduling scheme as an example to describe how to estimate efficiently these parameters via stochastic approximation techniques. Similar parameter estimation algorithms can be used for other scheduling schemes. Note that the authors in [46, 47] do not provide any algorithm to estimate their parameters, and the adaptive algorithm given in [51] cannot be implemented directly here because it involves equality constraints instead of the more general inequality constraints studied in Section 4.2.

To use a stochastic approximation algorithm to estimate $\vec{v}^*$, recall from Section 4.1.2 that $\vec{v}^*$ is chosen to satisfy the following condition: for any user i, if $v_i^* > \min_j v_j^*$, then $P\{Q^*(\vec{U}) = i\} = r_i$. Hence, we can write $\vec{v}^*$ as a root of the equation $f(\vec{v}) = 0$, where the ith component of $f(\vec{v})$ is given by

$$f_i(\vec{v}) = \left(v_i - \min_j(v_j)\right)\left(P\{Q(\vec{U}) = i\} - r_i\right), \qquad i = 1, \cdots, N.$$

Next, we use a stochastic approximation algorithm to generate a sequence of iterates $\vec{v}^1, \vec{v}^2, \cdots$ that represent estimates of $\vec{v}^*$. Each $\vec{v}^k$ defines a policy Q^k given by

$$Q^k(\vec{U}) = \operatorname*{argmax}_i(U_i + v_i^k).$$

To construct the stochastic approximation algorithm, we need an estimate g^k of $f(\vec{v}^k)$. Although we cannot obtain $f(\vec{v}^k)$ directly, we have a noisy observation of its components:

$$g_i^k = \left(v_i^k - \min_j(v_j^k)\right)\left(\mathbf{1}_{\{Q^k(\vec{U})=i\}} - r_i\right), \qquad i = 1, \cdots, N.$$

The observation error in this case is

$$e_i^k = g_i^k - f_i(\vec{v}^k) = \left(v_i^k - \min_j(v_j^k)\right)\left(\mathbf{1}_{\{Q^k(\vec{U})=i\}} - P\{Q^k(\vec{U}) = i\}\right),$$

and thus we have $E(e_i^k) = 0$. Hence, we can use a stochastic approximation algorithm of the form

$$v_i^{k+1} = v_i^k - \delta^k g_i^k,$$

where, e.g., $\delta^k = 1/k$.

When $v_i^k = \min_j v_j^k$, we also need to ensure that $P\{Q^k(\vec{U}) = i\} \geq r_i$. If $P\{Q^k(\vec{U}) = i\} < r_i$, then $\vec{v}^k$ is an infeasible parameter vector, which causes some fairness constraint to be violated. To ensure that $\{v_i^k\}$ converges to v_i^*, we should project $\vec{v}^k$ onto the feasible set of $\vec{v}$. However, because we do not have knowledge of the distribution of $\vec{U}$, it is very difficult to find the exact projection. Hence, we use the following intuitive algorithm as a projection.

It is easy to see that $P\{Q(\vec{U}) = i\}$ is an increasing function of v_i. Hence, if $v_i = \min_j v_j$ and $P\{Q(\vec{U}) = i\} < r_i$, then we increase the value of v_i to increase the value of $P\{Q(\vec{U}) = i\}$, as a projection to the feasible set. Although we do not know the value of $P\{Q(\vec{U}) = i\}$, we can estimate it by a moving average. Let p_i^k be the estimate of $P\{Q^k(\vec{U}) = i\}$. We update p_i^k in each time-slot by

$$p_i^k = (1-w)p_i^{k-1} + w\mathbf{1}_{\{Q^k(\vec{U})=i\}},$$

where w is a constant, indicating how fast p_i^k tracks $P\{Q^k(\vec{U}) = i\}$. If $p_i^k < r_i$ and $v_i^k = \min_j v_j^k$, then we update v_i^k as

$$v_i^{k+1} = v_i^k + \Delta,$$

where Δ is a positive constant. By doing this, we push $\vec{v}^k$ towards the feasible set of $\vec{v}$. We will show via simulations that this approach works well.

4.8 Simulation Results

In this section, we present numerical results from computer simulations of our scheduling schemes. For the purpose of comparison, we also simulate two special scheduling policies. The first is round-robin, a non-opportunistic scheduling policy that schedules (active) users following a predetermined order. This scheduling scheme serves as a benchmark of the system performance in order to measure how much gain the system can obtain using opportunistic scheduling policies. The second is a greedy scheduling scheme that always selects the user with the maximum performance at a generic time-slot to transmit. This greedy policy provides a tight upperbound on the system performance as explained earlier in Section 4.5, and is used here to indicate the tradeoff between the levels of QoS required by individual users and the overall system performance. In general, the looser the requirements, the better the overall performance.

We use the same cellular model as in Section 3.5. Figure 4.2 shows the forms of the performance values used by different users (there are 10 users in the system). The performance values of users 1, 5, and 8 are step-functions of SINR. The performance values of users 2, 6, and 9 are linear functions of SINR (in dB). Users 3, 4, 7, and 10 have performance values that are S-shape functions of SINR. Here we assume that users always have enough information to transmit when they are "on".

The 10 users are divided into three "distance" groups. Specifically, when a user becomes active, its distance from the base station is fixed, depending on which group it belongs to. Users 1–4 belong to the "far" group, i.e., when the user becomes active, its distance from the base station is $0.9R$, where R is the radius of the cell. Users 5–7 belong to the "middle" group; their starting distance from the base station is $0.5R$. Users 8–10 belong to the "near" group with a starting distance $0.2R$. When the user is active, it moves around in the cell freely and randomly. However, a user in a "near" group has a much higher chance to be close to the base station than a user in a "far" group. Hence, we can study how the distance from the base station effects users' performance in different scheduling schemes.

In the following paragraph, we describe the simulation procedure for the temporal fairness scheduling scheme. Other scheduling schemes follow the same simulation procedure, except for the details in steps 3 and 7. At the beginning of the simulation, we set the initial value of the parameter, i.e., $\vec{v}^1 = \vec{0}$.

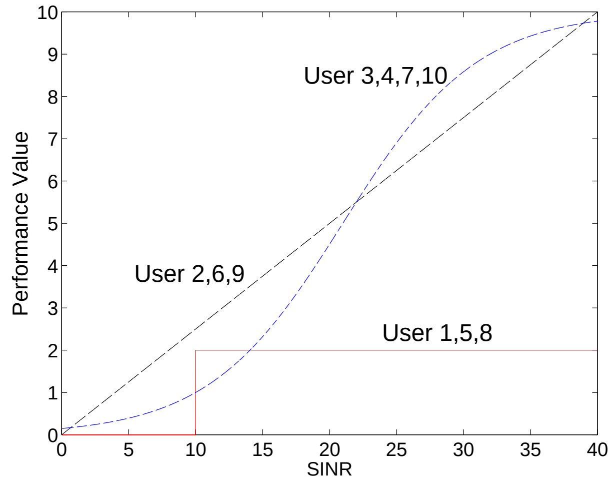

Fig. 4.2. Users' performance values as a function of SINR.

We maintain an ordered list of users in the system for the round-robin scheduling scheme. Let N be the number of active users. At each time-slot $k = 1, 2, \cdots$, the following steps are simulated:

1. If user i is active, we generate U_i^k. In our simulation, each user's performance value is a specific function of its SINR, as shown in Figure 4.2. Each user measures the received power level from the central base station, and the interference power level received from neighboring cells. We assume perfect measurement in all the simulations unless otherwise specified. Based on these measurements, the user calculates the SINR, and thus the corresponding performance value as a function of SINR.

2. Active users transmit their values of U_i^k to the base station through a signaling channel.

3. Based on the vector of performance values $\vec{U}^k$, the base station decides which user to schedule in the time-slot by

$$Q^k(\vec{U}^k) = \underset{i}{\operatorname{argmax}}(U_i^k + v_i^k).$$

4. If user $j = Q^k(\vec{U}^k)$ is the selected user, then the base station transmits to user j in the time-slot k. The system receives a performance "reward" equal to the performance value U_j^k.

5. In the round-robin scheduling scheme, we set J to be the index of the next active user in our ordered list of users, and let user J transmit. The system receives a performance "reward" equal to the performance value U_J^k of user J.

6. In the greedy scheduling scheme, we select the user I that has the maximum performance value, and let it transmit. The system receives a "reward" U_I^k.

7. The base station updates the parameters used in the scheduling policy for all active users, as described in Section 4.7:

$$v_i^{k+1} = v_i^k - \delta^k \left(v_i^k - \min_j(v_j^k) \right) \left(\mathbf{1}_{\{Q^k(\vec{U})=i\}} - r_i \right), \qquad i = 1, \cdots, N.$$

   We set $\delta^k = 0.01$ to track changes in the system because we are simulating a system that is not stationary in general. The larger the value of δ^k, the faster the estimated parameter tracks the actual value of the parameter (e.g., v_i^k tracks v_i^*), but at the same time the larger the fluctuation of the estimated parameter.

   Let P_j^k be the estimate of $P\{Q^k(\vec{U}^k) = j\}$; we update it according to

$$P_j^{k+1} = (1 - \epsilon_0)P_j^k + \epsilon_0 \mathbf{1}_{\{Q^k(\vec{U}^k)=j\}},$$

   where $\epsilon_0 = 0.001$ in the simulation. If $v_j^k = \min_i v_i^k$ and $P_j^k < r_j$, then

$$v_j^{k+1} = v_j^k + \Delta,$$

   where $\Delta = 0.02$ in the simulation. By doing this, we push v_i^ks to the feasible region.

The system performs the above procedure at every time-slot. Whenever the number of active users changes, the base station may need to update its QoS requirement, and the value of the parameters ($\vec{v}^k$) at that time is used as the initial value of the on-line parameter estimation procedure in the new system state.

4.8.1 Temporal Fairness Scheme

In this scheduling scheme, each user is entitled to a minimum portion of the resource. For our simulation, we set all users to have the same minimum resource requirement. Specifically, if N is the number of active users sharing the channel in the central cell, then each active user has a resource requirement $r_i = 1/(N + 3)$. In this case, $\sum_{i \in A} r_i = N/(N+3) < 1$, where A is the set of active users. Hence, the system has the freedom to assign a portion of the resource $[3/(N+3)]$ to some "fortunate" users (beyond their minimum requirements) to further improve the system performance, as discussed earlier.

Figures 4.3 and 4.4 show the results of our simulation experiments. In both figures, the x-axis represents the users' IDs. In Figure 4.3, the y-axis represents the portion of resource each user gets in the different scheduling policies. The first bar is the result of round-robin, where the resource is equally shared by all active users. Note that the amounts of resource consumed by different users may not be equal because different users are active at different times. The second bar shows the minimum requirements of users, while the third bar shows the portion of resource used by users in our temporal fairness scheduling scheme. The third bar is higher than the second bar for all the users, which indicates that our scheduling policy satisfies the fairness requirements of all users. Users 9 and 10 are the "fortunate" users in the system because they are most likely to be close to the base station and have large performance values. Thus, they get a much larger share of the resource than their minimum-requirements. The rightmost bars represent the greedy scheduling scheme that always chooses the user with the largest performance value to transmit. In this scheme, users 1, 2, 3, 4, 5, and 8 get very little or almost zero resource while users 9 and 10 have very large shares. As expected, the greedy algorithm is biased very heavily towards the users close to the base station with higher performance values (but, of course, violates the temporal fairness requirements).

Figure 4.4 shows the average performance obtained by users in different scheduling policies. The first bar represents round-robin, the second bar represents our scheduling policy, and the third bar is the greedy algorithm. Note that in the opportunistic scheduling scheme, all users except users 5 and 8 obtain higher average performance values than in the round-robin scheme. Moreover, users 1, 2, 3, 4, and 6 actually consume less resource while achieving better performance compared to round-robin, because users are more likely to be chosen to transmit while experiencing good channel conditions in the opportunistic scheduling scheme. To explain why users 5 and 8 do not have higher performance values, recall that both users 5 and 8 are inelastic users, whose performance values are step-functions of the SINR. Furthermore, user 5 belongs to the "middle" distance group, and user 8 is in the "near" group. Most of the time, they experience SINR values that are higher than the threshold in the step-function. Hence, there is little chance to improve their performance opportunistically. Compared with round-robin, these two users get smaller performance values because they consume less resource (recall that round-robin simply allocates the resource equally to all active users). We should clarify that these two users still obtain performance values that are greater than $r_i E(U_i)$, $i = 5, 8$, where r_i is the required time-fraction of user i, which is smaller than the fraction that user i gets in the round-robin scheme in this simulation.

The overall system performance is improved by 64% in our scheduling scheme while the greedy algorithm has a 111% performance improvement compared with round-robin. Of course, the greedy algorithm achieves this performance improvement by violating the QoS requirements of certain users.

We ignore the effects of fast multi-path fading in most of the simulations because it is not very clear how well fast fading can be tracked in real fields. However, we study the effect of fast Rayleigh fading in the following simulation. We adopt the simulation model for fast Rayleigh fading in [32], which amplitude modulates the inphase and quadrature components of a carrier with a low-pass filtered zero-mean Gaussian noise source. We consider two cases: First, fast fading can be accurately

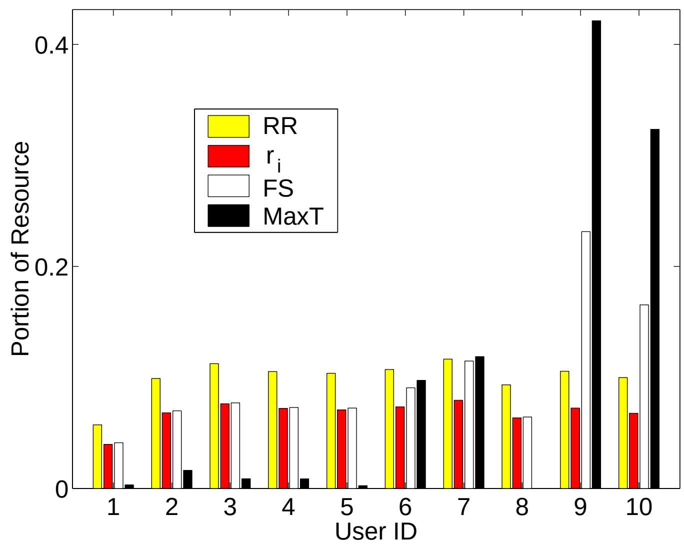

Fig. 4.3. Portion of resource shared by users in the temporal fairness scheduling simulation.

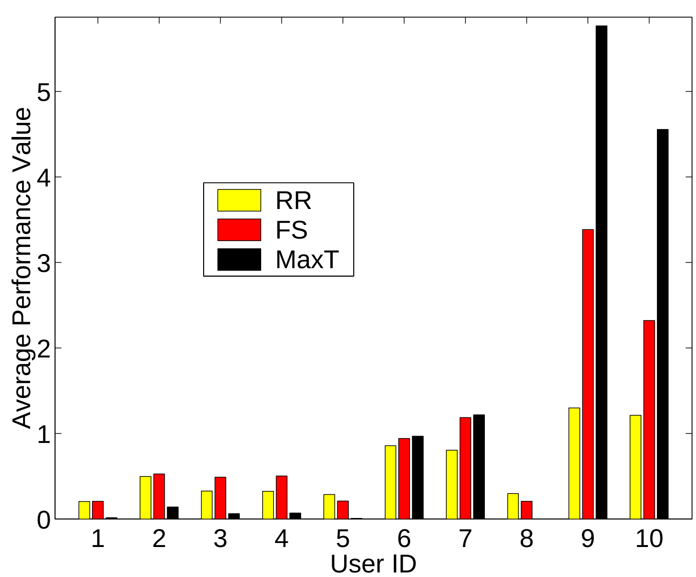

Fig. 4.4. Average performance value in the temporal fairness scheduling simulation.

estimated. Second, fast fading cannot be estimated, and thus there exists estimation error (of U_i) in the simulation. Note that in other simulations in the dissertation, fast fading is not simulated, and perfect estimation of U_i are assumed. Hence, besides studying the effect of fast fading, this simulation also provides a preliminary result on the robustness of our scheduling scheme and its implementation.

Figures 4.5 and 4.6 show the results of our simulation experiments in the fast fading environment. In both figures, the x-axis represents the users' IDs. In the legend, suffix "ff" represents the case that fast fading can be accurately estimated, and "ffn" means that fast fading is not estimated and thus acts as noise. In Figure 4.5, the y-axis represents the portion of resource each user gets in different scheduling policies. The first bar is the result of round-robin in the fast fading environment. The second bar shows the minimum requirements of users. The third and fourth bars show portion of resource used by users in our resource-based fairness scheduling policy in the cases where the fast fading is estimated and not estimated, respectively. In both cases, our scheduling policy satisfies the fairness requirement of all users. The rightmost two bars represent the greedy scheduling scheme with and without fast fading estimation, respectively.

Figure 4.6 shows the average performance obtained by users in different scheduling policies. The first bar represents round-robin in the fast fading environments. The second and third bar represent our scheduling policy with and without fast fading estimation, respectively. The last two bar are the results of the greedy algorithm, respectively. Compare these results with the simulation when fast fading is not considered, we can see that these figures are very similar. When fast fading exists and cannot be estimated, it introduces errors to the estimates of users' performance values. Hence, the system performance degrades to some extend in both our scheduling scheme and the greedy scheme. In the fast fading environment, the overall system performance is improved by 72% in our scheduling scheme when fast fading can be estimated accurately and by 67% when fast fading is not estimated compared with round-robin. For the greedy algorithm, the improvements are 119% and 117%, respectively. All of them are higher than the corresponding improvements in the previous simulation when fast fading is not simulated. The result is not surprising because the larger the variance of users' performance values, the higher the gain of opportunistic scheduling over non-opportunistic scheduling schemes in general. In summary, the simulation results show that the opportunistic scheduling policy can satisfy the fairness constraint with significant system performance gains, and our scheme is robust to estimation errors.

It is worth noting that the effect of fast fading needs further study. Our simulation model here is simplified. We assume fast fading is constant during a generic time-slot, and it only effects the SINR of users. In practice, fast fading can have substantial effects on error bit rates, and may result in retransmissions. Hence, tracking of fast fading and suitable channel coding schemes (such as redundancy incremental coding) should be further studied.

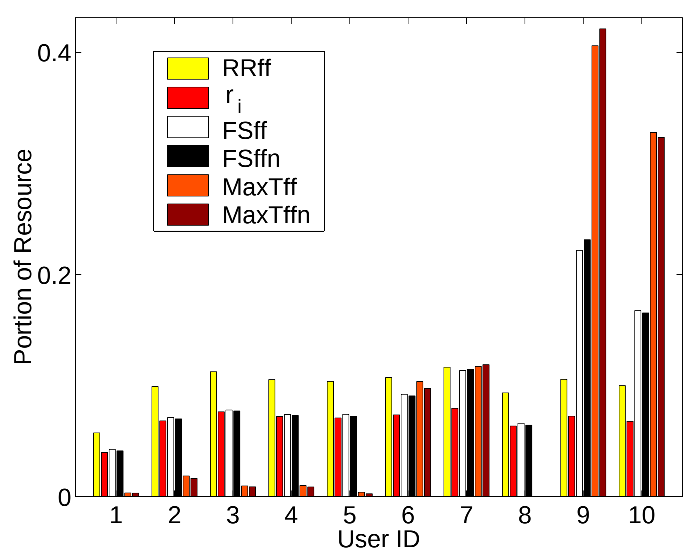

Fig. 4.5. Portion of resource shared by users for the resource-based fairness scheduling scheme in the fast fading simulation.

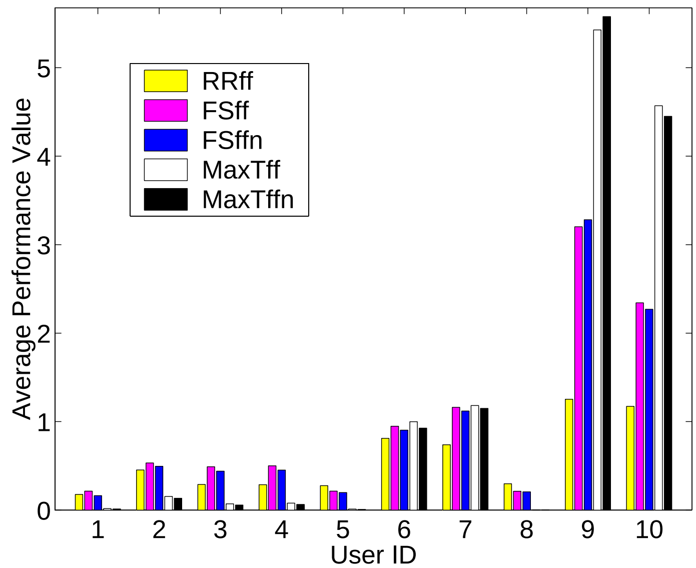

Fig. 4.6. Average performance value for the resource-based fairness scheduling scheme in the fast fading simulation.

4.8.2 Utilitarian Fairness Scheme

In this section, we show results for the utilitarian fairness scheduling scheme. We set the performance requirements of users as:

$$\vec{a} = [0.02 \;\; 0.08 \;\; 0.05 \;\; 0.05 \;\; 0.02 \;\; 0.1 \;\; 0.1 \;\; 0.02 \;\; 0.1 \;\; 0.1],$$

where $\sum_i a_i = 64\%$. In this setting, users 1, 5, and 8 have relatively low percentage requirements. Recall that these users have low and inelastic performance values. Therefore, to achieve the same amount of performance as other users, they require a larger portion of the resource. Hence, we assign these users low requirements to prevent them from using up too much resource in the system.

Figure 4.7 shows the average performance values of the users in different scheduling schemes. The x-axis represents the users' IDs, and the y-axis is the average performance value. In the figure, the first bar indicates the average performance values of users via the round-robin scheduling scheme. The second bar is the minimum performance value the user should get, which is calculated as:

$$a_i \frac{\sum_{k=1}^{T} U^k_{Q(\vec{U}^k)} \mathbf{1}_{\{\{\text{user } i \text{ is active}\}\}}}{\sum_{k=1}^{T} \mathbf{1}_{\{\{\text{user } i \text{ is active}\}\}}},$$

where T is the total number of time-slots simulated, and T =1,000,000. The denominator is the amount of time the user is active, whereas the numerator is the amount of system performance when the user is active. In other words, a user requires a share of the system performance only when it is active. The third bar indicates the performance value obtained from our scheduling policy. We can see in Figure 4.7 that all users obtain performance values larger than their minimum requirements. Moreover, most users achieve higher performance values than that of the round-robin scheme although there is no guarantee that the utilitarian scheme outperforms round-robin for each user. The rightmost bar is the result of the greedy algorithm. Compared with the round-robin scheme, our scheduling policy improves the overall system performance by 50% while the greedy scheduling has an improvement of 110%.

Figure 4.8 shows the average amount of resource consumed by different users. The first bar represents round-robin scheduling policy, the second bar represents our scheduling policy, and the third bar is the greedy scheduling policy. The greedy scheduling scheme allocates most resource to users in very good conditions (users 9 and 10). The utilitarian scheme, which also favors good users, allocates more resource (than the greedy algorithm) to other users to satisfy their fairness constraints.

4.8.3 Minimum-performance Guarantee Scheme

Next, we show simulation results for the opportunistic scheduling scheme with minimum-performance guarantees. First, we run the simulation for 1,000,000 time-slots using the round-robin scheduling policy, where the resource is equally distributed among all users, and active users are scheduled in a predetermined order. Thus we get an average performance value and use it as the minimum-performance requirement.

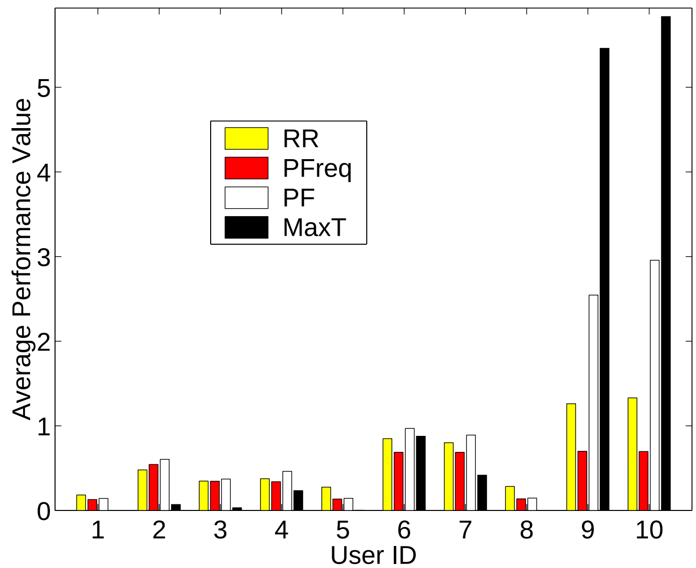

Fig. 4.7. Average performance value in the utilitarian fairness scheduling simulation.

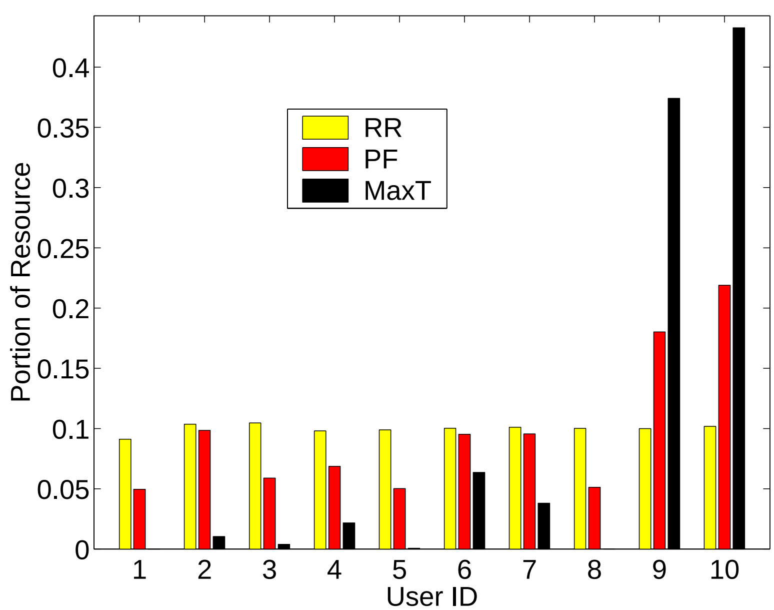

Fig. 4.8. Portion of resource shared by users in the utilitarian fairness scheduling simulation.

Then we run the simulation using the opportunistic scheduling policy, the round-robin policy, and the greedy scheduling policy.

Figure 4.9 shows the average performance values of users resulting from the different scheduling policies. The first bar indicates the average performance values using the round-robin scheduling policy, the second bar is the minimum-performance requirement of a user, the third bar indicates the result from our scheduling policy, and the rightmost bar is that of the greedy algorithm. Note that our scheduling policy outperforms the round-robin policy uniformly, which illustrates the "no-worse" guarantee discussed earlier in Section 4.3.2. Compared with round-robin, our scheduling policy improves the overall system performance by 51% while the greedy scheduling has an improvement of 109%. Similar to the previous simulation results, the greedy algorithm results in the highest overall performance value at the cost of the extreme unfairness among users.

Figure 4.10 shows the amount of resource consumed by each user in different scheduling policies. The first bar represents round-robin, the second bar represents our scheduling policy, and the third bar is the greedy scheduling scheme. As before, the greedy algorithm results in the most biased time-fraction allocation.

In summary, the simulations show that using our scheduling policies, the system can achieve significant performance gains while satisfying the QoS requirements. In the simulations with fairness constraints, there is no guarantee that every user performs better than using the round-robin policy. In the simulations with the minimum-performance requirement, we set the requirement to be the performance value obtained from round-robin; consequently, all users perform better using our policy than that in round-robin. In all the simulations, the greedy scheme, as expected, has the best performance at the cost of extreme unfairness among users, which indicates the possible tradeoff between users' QoS requirements and the system performance gain.

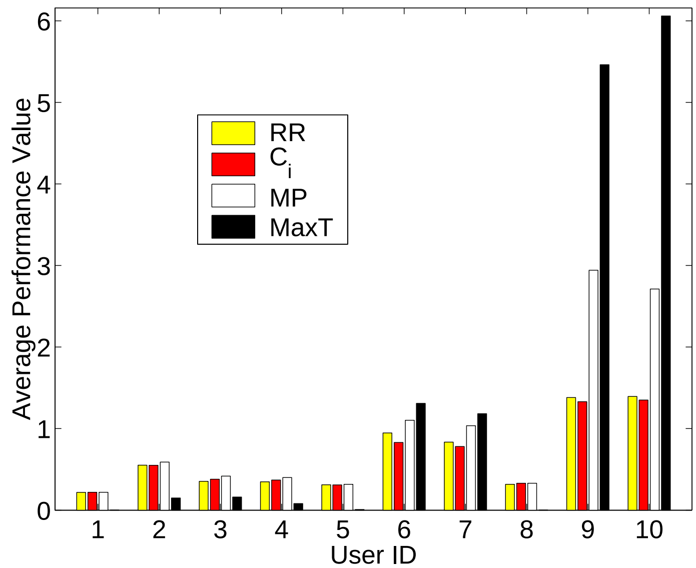

Fig. 4.9. Average performance value in the minimum-performance guarantee scheduling simulation.

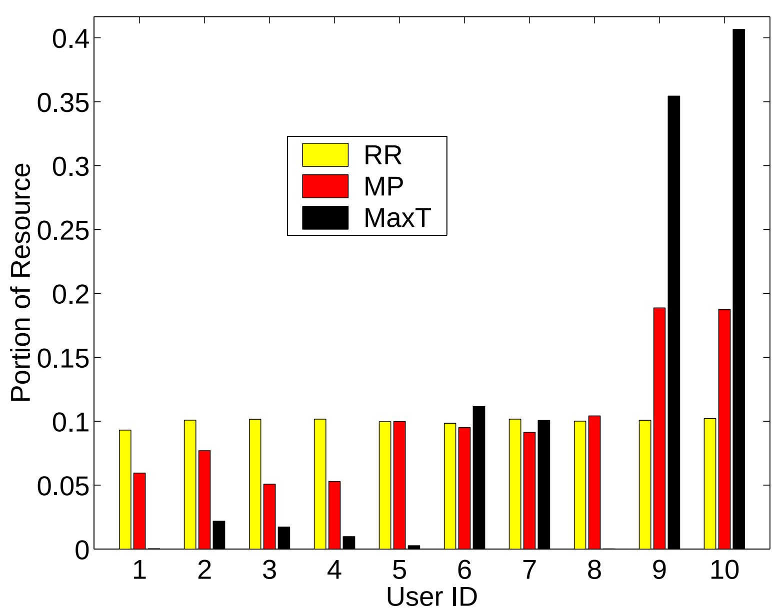

Fig. 4.10. Portion of resource shared by users in the minimum-performance guarantee scheduling simulation.

4.9 Conclusion

Opportunistic scheduling is a way to improve spectrum efficiency by exploiting time-varying channel conditions. In this chapter, we present a framework for opportunistic scheduling—to maximize the average system performance value by exploiting variations of the channel conditions while satisfying certain fairness/QoS constraints. The framework provides the flexibility to study a variety of opportunistic scheduling problems (many of the previous works by us and other researchers fit well into this unified framework). Using this framework, we have studied three scheduling schemes: to maximize the system performance with a temporal fairness requirement, a utilitarian fairness requirement, and a minimum-performance requirement for each user. We provide optimal solutions to each scheduling problem, and discuss their properties. Different scheduling schemes may be suitable for different application scenarios. We also study the asymptotic behavior of opportunistic scheduling schemes. Further, we extend our results to nonstationary policies under more general conditions. Lastly, we show via simulations that our opportunistic scheduling schemes result in substantial system gains while maintaining users' QoS requirements.

5. JOINT SCHEDULING AND POWER ALLOCATION

Wireless spectrum efficiency is becoming increasingly important with the growing demand for wide-band wireless services. Interference management is a crucial component of efficient spectrum utilization in wireless systems. Power allocation and opportunistic scheduling are effective mechanisms for interference management and efficient spectrum utilization.

In wireless communication systems, the received power represents signal strength to the desired receiver but interference to all other users. Power control is intended to provide each user an acceptable connection by eliminating unnecessary interference. The elegant work of Yates [8] abstracts the important properties of various power control algorithms and presents a unified treatment of power control. While power control is widely implemented in CDMA systems, such as IS-95 [9], it has also been shown to increase the call carrying capacity for channelized systems, such as TDMA/FDMA systems [10]. Furthermore, beyond the conventional concept of power control as a means to eliminating the "near-far" effect, power control is an effective resource management mechanism. It plays important role in interference management, channel-quality/service-quality provisioning, capacity management, and etc. [11, 10, 12, 13].

As explained in the previous chapters, opportunistic scheduling is a way to improve spectrum efficiency by exploiting the variation of wireless channels. An advantage of opportunistic scheduling is that it can be coupled with other resource management mechanisms to further improve network performance. In this chapter, we study the intercell-interference-alleviation problem using joint scheduling and power-allocation mechanisms. We study two different versions of this problem. In the first problem, the objective is to minimize the total transmission power, and thus interference to other cells, subject to a minimum-data-rate requirement for each user within the cell. In the second problem, the objective is to maximize the system net utility, which is defined as the value of the throughput minus the power cost, with the same minimum-data-rate constraints. In both problems, we have joint scheduling and power-allocation decisions.

5.1 System Model

As explained earlier, we consider a cell in a time-slotted system, and use a random variable α_i to represent the received SINR for user i at a generic time-slot given that the transmission power is 1. Let $\vec{\alpha} = \{\alpha_1, \cdots, \alpha_N\}$, where N is the number of users in the system. Basically, $\vec{\alpha}$ indicates the channel conditions of users at a generic time-slot. Let $f_i(c)$ be the required SINR for reliable transmission at data-rate c for user i, which is an increasing function of c, and $f_i(0) = 0$. Different users may have different forms of $f_i(c)$. Given α_i, which indicates the channel condition, the minimum required transmission power to support a data-rate c for user i is $f_i(c)/\alpha_i$. Let P_{max} be the maximum transmission power, which represents the restriction on the transmission power of a practical system. Let C_i denote the required average data-rate for user i.

There are two components in a joint scheduling and power-allocation scheme: a scheduling policy that decides which user to use the time-slot and a power allocation policy that decides the transmission power of the selected user (and thus its corresponding data-rate). Let Q be a scheduling policy; Q decides which user should transmit at a generic time-slot, given the channel conditions. In general, $Q(\vec{\alpha}) \in \{1, \cdots, N, \text{Null}\}$. If $Q(\vec{\alpha}) = i$, $i = 1, \ldots, N$, then user i is scheduled to transmit. If $Q(\vec{\alpha}) = \text{Null}$, then no user is scheduled to transmit. This may occur if all users experience relatively bad channel conditions. Let $p(\cdot)$ be the power allocation policy, $0 \leq p(\vec{\alpha}) \leq P_{max}$. If user i is selected to transmit and its transmission power is $p(\vec{\alpha})$, then $c_i(p) = f_i^{-1}(\alpha_i p)$ is its achievable data-rate. In summary, a policy for a joint scheduling and power-allocation scheme is given in the form of $\{Q, p\}$.

5.2 Minimizing Transmission Power

First, we study the problem where the objective is to minimize the overall transmission power while maintaining the required data-rate for each user within the cell. By minimizing the overall transmission power, we decrease the intercell interference without communications among base stations. To achieve this goal, we need to decide which user should be scheduled at a generic time-slot and what should be its transmission power. Let $P(Q,p)$ be the overall transmission power of the system under policy $\{Q,p\}$:

$$P(Q,p) = \sum_{i=1}^{N} E\left(p(\vec{\alpha})\mathbf{1}_{\{Q(\vec{\alpha})=i\}}\right).$$

The problem that we are interested in can be formally stated as:

$$\begin{aligned} \underset{Q,p}{\text{minimize}} \quad & P(Q,p) \\ \text{subject to} \quad & 0 \le p(\vec{\alpha}) \le P_{max}, \\ & E\left(c_i(p)\mathbf{1}_{\{Q(\vec{\alpha})=i\}}\right) = C_i, \quad i = 1, \cdots, N. \end{aligned} \tag{5.1}$$

Our objective is to minimize the overall transmission power $P(Q,p)$ under two sets of constraints. The first constraint, $0 \le p(\vec{\alpha}) \le P_{max}$, indicates the maximum-transmission-power restriction of the system. The second constraint, $E(c_i(p)\mathbf{1}_{\{Q(\vec{\alpha})=i\}}) = C_i$, is the minimum-data-rate constraint, where $E(c_i(p)\mathbf{1}_{\{Q(\vec{\alpha})=i\}})$ is the average data-rate of user i given $\{Q,p\}$. Note that we could have written the second constraint in the more general (inequality) form: $E(c_i(p)\mathbf{1}_{\{Q(\vec{\alpha})=i\}}) \ge C_i$. However, because our objective is to minimize the transmission power, a solution to (5.1) is certainly a solution to the problem with the more general inequality constraints. Hence, without loss of generality, we study the problem with the equality constraints, as defined in (5.1).

We next present our solution to the joint scheduling and power-allocation problem defined in (5.1). Let

$$L(\vec{\lambda}) = \sum_{i=1}^{N} E\left(p(\vec{\alpha})\mathbf{1}_{\{Q(\vec{\alpha})=i\}}\right) - \sum_{i=1}^{N} \lambda_i \left(E\left(c_i(p)\mathbf{1}_{\{Q(\vec{\alpha})=i\}}\right) - C_i\right)$$

$$= \quad E\left(\sum_{i=1}^{N}\left(p(\vec{\alpha})-\lambda_i c_i(p)\right)\mathbf{1}_{\{Q(\vec{\alpha})=i\}}\right)+\sum_{i=1}^{N}\lambda_i C_i.$$

We define

$$\begin{aligned} l_i(\vec{\lambda},\vec{\alpha},p) &= p(\vec{\alpha})-\lambda_i c(p) \\ p_i^*(\vec{\lambda},\vec{\alpha}) &= \operatorname*{argmin}_{0\le p\le P_{max}} l_i(\vec{\lambda},\vec{\alpha},p) \\ l_i^*(\vec{\lambda},\vec{\alpha}) &= l_i(\vec{\lambda},\vec{\alpha},p_i^*(\vec{\lambda},\vec{\alpha})). \end{aligned}$$

Note that we have $l_i^*(\vec{\lambda},\vec{\alpha})\le 0$ because $l_i(\vec{\lambda},\vec{\alpha},0)=0$.

Proposition 10 *Suppose there exists* $\vec{\lambda}^*$ *such that*

$$E\left(c_i(p^*)\mathbf{1}_{\{Q^*(\vec{\alpha})=i\}}\right)=C_i, \quad i=1,\cdots,N,$$

where $Q^*(\vec{\alpha})$ *is defined as*

$$Q^*(\vec{\alpha})=\operatorname*{argmin}_i l_i^*(\vec{\lambda}^*,\vec{\alpha}). \tag{5.2}$$

Then, $\{Q^*,p^*\}$ *is an optimal solution to the problem defined in (5.1).*

The above proposition is valid for all f_is that are increasing functions with $f_i(0)=0$. Further, if f_i is a strictly convex function, then p^* has a closed-form expression:

$$p^*(\vec{\lambda},\vec{\alpha})=\begin{cases} 0 & \text{if } f_i'(0)>\lambda_i\alpha_i \\ \frac{f_i(f_i'^{-1}(\lambda_i\alpha_i))}{\alpha_i} & \text{if } f_i'(0)\le\lambda_i\alpha_i\le f_i'(C_i) \\ P_{max} & \text{if } f_i'(C_i)<\lambda_i\alpha_i \end{cases},$$

where $C_i=f_i^{-1}(\alpha_i P_{max})$ is the maximum data-rate of user i given P_{max} and the channel condition.

From Proposition 10, we observe that a user is chosen to transmit when it is a "relatively-best" user. User i is "relatively-best" if $l_i^*(\vec{\lambda},\vec{\alpha})\le \min_j l_j^*(\vec{\lambda},\vec{\alpha})$ (and hence has the same form as the opportunistic scheduler in Chapter 4). Moreover, the transmission power of the selected user is the power that minimizes $l_i(\vec{\lambda},\vec{\alpha},p)$. We can think of λ_i as the *unit reward* (in terms of power/data-rate) to compensate power consumption. It controls the value of transmission power, and in turn the data-rate. The fact that $l_i^*(\vec{\lambda}, \vec{\alpha}) \leq 0$ indicates that the transmission power $(p^*(\vec{\lambda}, \vec{\alpha}))$ should be no greater than the reward $(\lambda_i c_i(p^*))$ of the user for transmitting at data-rate $c_i(p^*)$. Also note that the resulting data-rate of a user is an increasing function of its unit reward λ_i. This property enables us to obtain $\vec{\lambda}^*$ iteratively in the implementation.

Proof of Proposition 10: Suppose a policy $\{Q, p\}$ satisfies the maximum transmission power constraint and the data-rate constraint. We will show that

$$P(Q, p) \geq P(Q^*, p^*).$$

Because $\{Q, p\}$ satisfies the rate constraint, we have

$$\begin{aligned} P(Q, p) &= P(Q, p) - \sum_{i=1}^{N} \lambda_i^* \left(E(c_i(p))\mathbf{1}_{\{Q(\vec{\alpha})=i\}} - C_i\right) \\ &= E\left(\sum_{i=1}^{N} (p(\vec{\alpha}) - \lambda_i^* c_i(p))\, \mathbf{1}_{\{Q(\vec{\alpha})=i\}}\right) + \sum_{i=1}^{N} \lambda_i^* C_i. \end{aligned}$$

Further,

$$l_i(\vec{\lambda}^*, \vec{\alpha}, p_i^*(\vec{\lambda}^*, \vec{\alpha})) \leq l_i(\vec{\lambda}^*, \vec{\alpha}, p), \quad 0 \leq p \leq P_{max},$$

by the definition of $p_i^*(\vec{\lambda}^*, \vec{\alpha})$. If $Q \neq$ Null, then

$$\begin{aligned} &\sum_{i=1}^{N} (p(\vec{\alpha}) - \lambda_i^* c_i(p))\, \mathbf{1}_{\{Q(\vec{\alpha})=i\}} \\ \geq\ &\sum_{i=1}^{N} l_i(\vec{\lambda}^*, \vec{\alpha}, p^*(\vec{\lambda}^*, \vec{\alpha}))\mathbf{1}_{\{Q(\vec{\alpha})=i\}} \\ \geq\ &\sum_{i=1}^{N} l_i(\vec{\lambda}^*, \vec{\alpha}, p^*(\vec{\lambda}^*, \vec{\alpha}))\mathbf{1}_{\{Q^*(\vec{\alpha})=i\}}, \end{aligned}$$

where the last inequality is because of the definition of Q^*. If $Q =$ Null, because $l_i(\vec{\lambda}^*, \vec{\alpha}, p^*(\vec{\lambda}^*, \vec{\alpha})) \leq 0$, we have

$$\begin{aligned} &\sum_{i=1}^{N} (p(\vec{\alpha}) - \lambda_i^* c_i(p))\, \mathbf{1}_{\{Q(\vec{\alpha})=i\}} \\ &= 0 \\ &\geq \sum_{i=1}^{N} l_i(\vec{\lambda}^*, \vec{\alpha}, p^*(\vec{\lambda}^*, \vec{\alpha}))\mathbf{1}_{\{Q^*(\vec{\alpha})=i\}}. \end{aligned}$$

Hence,

$$\begin{aligned} P(Q,p) &\geq E\left(\sum_{i=1}^{N} l_i(\vec{\lambda}^*, \vec{\alpha}, p^*(\vec{\lambda}^*, \vec{\alpha}))\mathbf{1}_{\{Q^*(\vec{\alpha})=i\}}\right) + \sum_{i=1}^{N} \lambda_i^* C_i \\ &= P(Q^*, p^*). \end{aligned}$$

$\square$

5.3 Maximizing Net Utility

Now we study a joint scheduling and power-allocation problem in a different scenario. In Section 4.3, we consider a scheduling problem that maximizes the system throughput subject to each user's minimum data-rate requirement. Because a user's throughput is an increasing function of its transmission power and the objective is to maximize the throughput, it is obvious that the base station should always transmit with its maximum power. However, because transmission power causes interference to other cells in wireless systems, we need to take the power consumption into account as well. To achieve this goal, we introduce the notion of "net utility," which is defined as the difference between the value of the throughput and the cost of the power consumption [13]. Note that in the context of downlink transmissions, the transmission power itself is not a concern because base stations usually have adequate power supplies. Hence, the power cost here actually refers to the interference cost of the transmission power, which is different from [13]. Let $g_i(p)$ be the power cost of user i; then, $c_i(p) - g_i(p)$ is defined as the net utility of user i. As before, let Q be the scheduling policy and p be the power policy. Let $T_i(Q,p)$ be the average net utility of user i given the joint policy $\{Q,p\}$:

$$\begin{aligned} T_i(Q,p) &= E\left((c_i(p) - g_i(p))\mathbf{1}_{\{Q(\vec{\alpha})=i\}}\right) \\ T(Q,p) &= \sum_{i=1}^{N} T_i(Q,p). \end{aligned}$$

We formulate the problem as:

$$\underset{Q,p}{\text{maximize}} \qquad T(Q,p)$$

$$\text{subject to} \quad 0 \le p(\vec{\alpha}) \le P_{max}$$
$$E\left(c_i(p)\mathbf{1}_{\{Q(\vec{\alpha})=i\}}\right) \ge C_i, \quad i = 1, \cdots, N \tag{5.3}$$

In other words, the objective is to maximize the net utility given the maximum power constraint and data-rate requirement constraints. We define

$$b_i(\vec{\eta}, \vec{\alpha}, p) = c_i(p)\eta_i - g_i(p)$$
$$p^*(\vec{\eta}, \vec{\alpha}) = \max_{0 \le p \le P_{max}} b_i(\vec{\eta}, \vec{\alpha}, p)$$
$$b_i^*(\vec{\eta}, \vec{\alpha}) = b_i(\vec{\eta}, \vec{\alpha}, p^*(\vec{\eta}, \vec{\alpha})).$$

Let $Q^*(\vec{\alpha})$ be defined as

$$Q^*(\vec{\alpha}) = \operatorname*{argmax}_i b_i^*(\vec{\eta}^*, \vec{\alpha}), \tag{5.4}$$

where $\vec{\eta}^*$ satisfies:

1. $\min_{i=1,\cdots,N} \eta_i^* = 1$
2. $T_i(Q^*, p^*) \ge C_i$ for all i
3. For all i, if $T_i(Q^*, p^*) > C_i$, then $\eta_i^* = 1$.

Proposition 11 *$\{Q^*, p^*\}$ is an optimal solution to the problem defined in (5.3).*

Proof: Suppose $\{Q, p\}$ satisfies the maximum transmission power constraint and the data-rate constraint. We will show that $T(Q, p) \le T(Q^*, p^*)$. We have

$$\begin{aligned} T(Q,p) &\le T(Q,p) + \sum_{i=1}^{N} (\eta_i^* - 1)(T_i(Q,p) - C_i) \\ &= \sum_{i=1}^{N} E\left((c_i(p)\eta_i^* - g_i(p))\mathbf{1}_{\{Q(\vec{\alpha})=i\}}\right) - \sum_{i=1}^{N} (\eta_i^* - 1)C_i. \end{aligned}$$

Further,

$$\begin{aligned} &\sum_{i=1}^{N} (c_i(p)\eta_i^* - g_i(p))\mathbf{1}_{\{Q(\vec{\alpha})=i\}} \\ \le &\sum_{i=1}^{N} (c_i(p^*)\eta_i^* - g_i(p^*))\mathbf{1}_{\{Q(\vec{\alpha})=i\}} \\ \le &\sum_{i=1}^{N} (c_i(p^*)\eta_i^* - g_i(p^*))\mathbf{1}_{\{Q^*(\vec{\alpha})=i\}}. \end{aligned}$$

The first inequality is due to the definition of p^*, and the second to the definition of Q^*. Hence,

$$\begin{aligned} T(Q,p) &\leq \sum_{i=1}^{N} E\left((c_i(p^*)\eta_i^* - g_i(p^*))\mathbf{1}_{\{Q^*(\vec{\alpha})=i\}}\right) - \sum_{i=1}^{N}(\eta_i^* - 1)C_i \\ &= \sum_{i=1}^{N} T_i(Q^*, p^*) + \sum_{i=1}^{N}(\eta_i^* - 1)\left(T_i(Q^*, p^*) - C_i\right) \\ &= T(Q^*, p^*), \end{aligned}$$

which completes the proof. $\square$

Note that the solutions of the two problems have certain similarities. Both policies choose the "relatively-best" user to transmit and the optimal transmission power maximizes/minimizes b_i/l_i. Here, a user is relatively-best if $b_i^*(\vec{\eta}^*, \vec{\alpha}) \geq \max_j b_j^*(\vec{\eta}^*, \vec{\alpha})$.

The problem studied in Section 4.3 can be considered as a special case of the problem defined in (5.3) with $g_i(p) \equiv 0$. In other words, if we do not penalize transmission power at all, then the base station always transmits with its maximum power P_{max}. In this case, the joint scheduling and power-allocation degenerates to a pure scheduling problem.

5.4 Discussion

So far, we have studied three scheduling problems with the minimum-data-rate constraints. In the first problem formulation, defined as (4.5) in Section 4.3, the objective is to maximize the total performance value of a cell given the minimum-data-rate constraints. The second problem, defined in (5.1), is to minimize the transmission power, and the third, defined in (5.3), is to maximize the net utility, both under the same constraint. These problem formulations target for different scenarios. In the first problem, there is absolutely no collaboration among base stations. Each base station transmits with its maximum power, and thus generates maximum interference to other base stations. In the second and third problems, base stations are more "considerate": they take into account the negative effect of their transmission powers on other cells (i.e., intercell interference). In the first and second problems, the system has different objectives. It is clear that using the second problem formulation, the system can admit more users while satisfying their minimum data-rate requirements or each user can have less degradation probability. On the other hand, a user may get higher throughput (better than what it asks for) in the first formulation. The third problem formulation targets to integrate the merits of the first two. If the power cost is set high, then the third problem may degenerate to the second one. If the power cost is set low, then the result may be the same as the first one. By setting a suitable cost function, the base station can balance the tradeoff between the throughputs of its own users and considerations for other cells, and thus improve the overall system performance (e.g., throughputs/dropping probability in all cells). Hence, a challenging problem is to find a cost function that reflects how the transmission power of one cell affects the capacity of other cells (on average).

Because the transmission power of one base station causes interference to users in other cells, one potential effect of the power allocation scheme is to induce the fluctuation of channels in other cells, and thus may improve the scheduling gain, especially in environments with little scatters and/or slow fading. Yet, there are many questions to be answered in this context. The first question is how to ensure the fluctuation in an appropriate time-scale such that we can increase the opportunistic scheduling gain without experiencing significant delay. Second, we should study that under what conditions which scheme performs better in terms of the number of users admitted, the probability of failure, and the overall throughput. Further research also includes the study of the system behavior (e.g., convergence of the transmission power) when all cells implement the same scheduling and power-allocation algorithms.

5.5 Numerical Results

We use the same cellular model as in Section 3.5. Recall that $f(c)$ is the required SINR for reliable transmission at data-rate c. In practice, f is usually not a continuous function because the system only supports discrete data-rates by adapting different

Rate (kbps)	10	20	30	40	50	60	70	80
SINR (dB)	5	10	15	20	25	30	35	40

Table 5.1
Achievable data-rate vs. SINR

coding rates and modulation schemes. Hence, in the simulation, we assume that there are ten discrete data-rates available and the corresponding SINR requirements are listed in Table 5.1. The data is similar to the result presented in [66].

In the following, we show simulation results for the joint scheduling and power-allocation scheme that minimizes the overall transmission power with minimum-data-rate guarantees. First, we set $\vec{\lambda}^0$. Then the system performs the following procedure at every time-slot. Basically, the users measure their channel conditions, and then send the information to the base station ($\vec{\alpha}^k$). The base station decides the user to be scheduled and its transmission power using the joint scheduling and power-allocation policy by substituting $\vec{\lambda}^k$ into (5.2). Then the base station updates the parameter used in the scheduling policy by

$$\lambda_i^{k+1} = \lambda_i^k + a^k(C_i - \bar{R}_i^k),$$

where a_k is the step size ($a_k = 0.001$), $\bar{R}_i^k$ is the estimated average data-rate of user i at time k, C_i is the required data-rate, and $C_i = 3$ (kbps).

For the purpose of comparison, we also simulate a round-robin policy. In the round-robin scheduling policy, active users follow a predetermined order. When a user's turn comes, if its average transmission rate is lower than its required rate, then we let the user transmit. Its transmission power is set to be the minimum power required to support the highest data-rate achievable to the user given its channel condition. If the use does not qualify to transmit, then we go to the next active user until we find a user to transmit in the current time-slot or all active users have been exhausted.

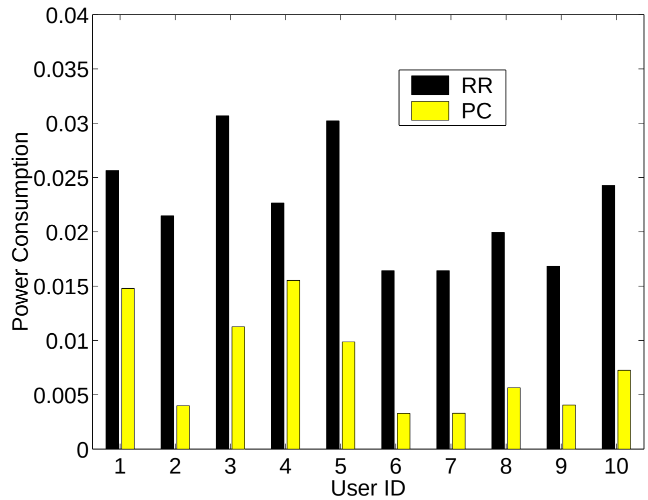

Fig. 5.1. The average transmission power of round-robin and our policy.

We compare the average transmission power of our policy with that of round-robin. The simulation is run for 100,000 time-slots. In the simulation, both policies can maintain the required data-rates of users. Figure 5.1 shows the average transmission power of each user using our policy and using round-robin. The x-axis is the user's ID. The y-axis is the average power consumption, which is the amount of transmission power consumed by a user divided by its throughput. The first bar is the average power consumption of a user in the round-robin scheme and the second bar is that in our scheme. Note that the power consumption of each user in our scheme is unanimously lower than that in round-robin. The total transmission power is only 32% of that of round-robin.

5.6 Conclusions

Opportunistic scheduling exploits time-varying channel conditions to improve spectrum efficiency, providing an additional degree of freedom to the system. Its merits also lie in its ability to work in conjunction with other resource management mechanisms. In this chapter, we study joint scheduling and power-allocation schemes to alleviate intercell interference. We study two problems with different but related objectives. The first objective is to minimize the total transmission power and the second objective is to maximize the net utility, both under minimum-data-rate constraints. We provide optimal solutions to the studied problems and use simulation results to evaluate the power savings compared with a round-robin scheme.

6. CONCLUSIONS AND FUTURE WORK

In this chapter, we first conclude our work on opportunistic scheduling, and discuss the advantages and disadvantages of opportunistic scheduling schemes. We then present some open problems in the area.

6.1 Conclusions

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. Opportunistic scheduling exploits the variation of channel conditions to improve spectrum efficiency. In this thesis, we have studied opportunistic scheduling for time-slotted systems. Such systems include TDMA and certain CDMA systems. First, the time-slot in the system model is a natural match for the time-slot in a TDMA system. Further, research shows that, to achieve high-data capacity, data users should transmit in a time-multiplexed mode instead of transmitting simultaneously as in traditional CDMA systems [56, 57, 25]. Hence, the time-slotted system is also appropriate for CDMA systems used for high-rate data communications.

To improve spectrum efficiency, intuitively, we want to assign resource to users experiencing "good" channel conditions. At the same time, it is also desirable to provide some form of fairness or QoS guarantees. Otherwise, the system performance can be trivially optimized by, for example, letting a user with the highest performance value to transmit. This may prevent "poor" users (in terms of either channel conditions or money) from accessing the network resource, and thus compromises the desirable feature of wireless networks: to provide "anytime," "anywhere" accessibility. In the dissertation, we study opportunistic scheduling schemes under a unified framework.

Each user's channel-condition/performance-value is modeled by a stochastic process, reflecting the time-varying performance that results from randomly-varying channel conditions. The objective is to maximize the system performance under certain QoS constraints. The framework provides the flexibility to study a variety of opportunistic scheduling problems (many previous investigates fit under this unified framework). Using this framework, we have studied three scheduling schemes: to maximize the system performance with a temporal fairness requirement, a utilitarian fairness requirement, and a minimum-performance requirement for each user. We find optimal solutions for these scheduling problems. An attractive feature of these optimal solutions is that they are given in a simple parametric form, hence lending themselves to on-line implementation. We provide algorithms to estimate the parameters in these optimal solutions and describe implementational procedures for each solution.

The schemes studied in the dissertation have different properties. In wireline networks, when a certain amount of resource is assigned to a user, it is equivalent to granting the user a certain amount of throughput/performance value. However, the situation is different in wireless networks, where the amount of resource and the performance value are not directly related (though closely correlated). Hence, we study both temporal and utilitarian fairness requirement in the dissertation. While fairness criteria provide users a relative measure of performance, we also study an absolute measure of performance, i.e., "minimum-data-rate". Although providing users with a minimum-data-rate requirement is desirable to some extent, it suffers from feasibility problems. First, it is not an easy task to check whether users' requirements can be met at a give time. Second, it is even harder to meet the requirements through the users' lifetime due to users' mobility and system capacity variation.

An advantage of opportunistic scheduling is that it can be coupled with other resource management mechanisms to further increase network performance. In Chapter 5, we study joint scheduling and power-allocation mechanisms for intercell-interference alleviation. Interference management is crucial for spectrum efficiency because interference ultimately limits the system capacity. Power control is a traditional interference management mechanism, and opportunistic scheduling exploits time-domain diversity. We combine the merits of the two and study different versions of interference management problems. In the first problem, the objective is to minimize the total transmission power, and thus interference to other cells, subject to a minimum-data-rate requirement for each user within the cell. In the second problem, the objective is to maximize the system net utility, which is defined as the difference between the value of the throughput and the power cost, with the same minimum-data-rate constraints. In both problems, we have joint scheduling and power-allocation decisions. These schemes are more considerate in the sense that they take into account the negative result of their transmission powers to other cells.

In summary, opportunistic scheduling exploits 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. The schemes studied in the dissertation are relatively easy to implement, and robust to system dynamics, such as the mobility of users and the status changes of users. 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.

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

- 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. In Chapter 3, we have some preliminary results to show that our scheme is robust to estimation errors. However, we have not taken into account the actual characteristics of the estimation procedure in this study. We recommend a rigorous study on this problem, especially in the case of fast fading.

- Opportunistic scheduling exploits the fluctuation of channel condition, and thus scheduling gain 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 most 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.

6.2 Future Work

Resource allocation and scheduling schemes are important in wireless networks, especially to provide high-rate data and seamless service. There are many interesting problems yet to be resolved in this area.

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 [46, 47, 48, 49, 50], where queueing status is a part of scheduling decisions, and [64], which is a heuristic extension of the GPS.

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 several approaches, e.g., [67, 68, 69, 70, 71]. 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 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, online model-based solutions, and heuristic algorithms [53]. 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 can be expected to hold for wireless networks.

Our opportunistic scheduling framework 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 (ant time) that are relatively good. An example of such systems is OFDM systems. OFDM (orthogonal frequency division multiplexing) is a promising transmission (modulation) technique to combat ISI over multipath fading channels and provide efficient frequency utilization [17, 21]. Resource management in OFDM systems has been studied in both single-user systems [72, 73, 74, 75, 76, 77, 78] and multi-user systems [79, 80, 81, 82, 83, 84, 85]. However, none of the these works exploit the variations of channel conditions in time. Since an OFDM symbol usually lasts $100\mu s$ to $1000\mu s$, for data service without strict delay requirements, we can expect to reap performance benefits by exploiting time-domain diversity too. The performance measure in such systems may include bit-rate, power consumption, and bit-error rate. The objective is to improve the system performance by exploiting both frequency-domain and time-domain diversity (with or without fairness/QoS constraints). 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.

Opportunistic scheduling exploits the channel fluctuations of users. Hence, the larger the channel fluctuation, the higher the scheduling gain. Then a natural question is what to do in environments with little scattering and/or slow fading. In [30], the authors use multiple transmission antennas to "induce" fluctuations of the received SINR of users, and thus exploit multi-user diversity. The merit of the work is to find the optimal beamforming without feedback information on each antenna (phase and amplitude). This scheme does not "induce" physical fluctuations in channels. For downlink transmission, the transmission power of one base station causes interference to users in other cells. Hence, using a joint power control and scheduling scheme, we can actually induce fluctuation physically by changing the transmission power at different base stations. The problems to be studied include the system behavior (e.g., convergence), time-scale of fluctuation, and overall capacity, etc., as discussed in Section 5.4.

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_is are i.i.d. with exponential distribution, then the scheduling gain is $O(\log(n))$. On the other hand, the signaling cost per user remains the same. Hence, it is an question of practical importance to decide the number of users sharing a same channel.

The opportunistic scheduling problems described in the previous chapters 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.

LIST OF REFERENCES

[1] S. Verdu. Wireless bandwidth in the making. *IEEE Communications Magazine*, 38(7):53–58, July 2000.

[2] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423 and 623–656, July and October 1948.

[3] J. Zhang. *Design and performance analysis of power-controlled CDMA wireless networks with linear receivers and antenna arrays.* PhD thesis, Purdue University, 2000.

[4] P. Lehne and M. Petterson. An overview of smart antenna technology for mobile communications systems. *IEEE Communications Surveys*, 2(4):2–13, 1999.

[5] A. Boukalov and S. Haggman. System aspects of smart-antenna technology in cellular wireless communications-an overview. *IEEE Transactions on Microwave Theory and Techniques*, 48(6):919 –929, 2000.

[6] G. J. Foschini. Layered space-time architecture for wireless communication in a fading environment when using multi-element antennas. *Bell Labs Technical Journal*, pages 41–59, Autumn 1996.

[7] C. Berrou, A. Glavieux, and P. Thitimajshima. Near Shannon limit error-correcting coding and decoding: Turbo codes. In *Proceedings of 1993 IEEE International Conference on Communications*, Geneva, Switzerland, May 1993.

[8] R. Yates. A framework for uplink power control in cellular radio systems. *IEEE Journal on Selected Areas in Communications*, 43(7):1341 –1347, 1995.

[9] A. Viterbi. *CDMA: Principles of Spread Spectrum Communication.* Addison-Wesley, 1995.

[10] J. Zander. Performance of optimum transmitter power control in cellular radio systems. *IEEE Transactions on Vehicular Technology*, 41(1):51–62, February 1992.

[11] J. Zander. Trends in resource management future wireless networks. In *Future Telecommunications Forum, FTF 99*, Beijing, China, December 1999.

[12] F. Berggren, S. Kim, R. Jantti, and J. Zander. Joint power control and intracell scheduling of DS-CDMA nonreal time data. *IEEE Journal on Selected Areas in Communications*, 19(10):1860 –1870, October 2001.

[13] M. Xiao, N. Shroff, and E. Chong. Utility-based power control in cellular wireless systems. In *Proceedings of IEEE INFOCOM 2001*, volume 1, 2001.

[14] I. Katzela and M. Naghshineh. Channel assignment schemes for cellular mobile telecommunication systems: A comprehensive survey. *IEEE Personal Communications Magazine*, 3(3):10–31, June 1996.

[15] D. A. Levine, I. F. Akyildiz, and M. Naghshineh. A resource estimation and call admission algorithms for wireless multimedia networks using the shadow cluster concept. *IEEE/ACM Transactions on Networking*, 5(1):1–13, February 1997.

[16] I. F. Akyildiz, J. McNair, L. C. Martorell, R. Puigjaner, and Y. Yesha. Medium access control protocols for multimedia traffic in wireless networks. *IEEE Network Magazine*, 13(4):39–47, July/August 1999.

[17] L. Cimini, J. Chuang, and N. Sollenberger. Advanced cellular internet service (ACIS). *IEEE Communications Magazine*, 36(10):150–159, October 1998.

[18] B. Wu and Q. Wang. Maximization of the channel utilization in wireless heterogeneous multiaccess networks. *IEEE Transactions on Vehicular Technology*, 46(2):437–444, 1997.

[19] J. Chuang and N. Sollenberger. Beyond 3G: wideband wireless data access based on OFDM and dynamic packet assignment. *IEEE Communications Magazine*, 38(7):78–87, July 2000.

[20] H. Rohling, T. May, K. Bruninghaus, and R. Grunheid. Broad-band OFDM radio transmission for multimedia applications. *Proceedings of the IEEE*, 87(10):1778–1789, 1999.

[21] Flarion Technologies. *Flarion business and technology overview.* http://www.flarion.com/technology/tech_papers.html.

[22] Flarion Technologies. *Mobilizing the high-speed Internet.* http://www.flarion.com/technology/tech_papers.html.

[23] S. Nanda, K. Balachandran, and S. Kumar. Adaptation techniques in wireless packet data services. *IEEE Communications Magazine*, 38(1):54–64, January 2000.

[24] A. Goldsmith and S. G. Chua. Adaptive coded modulation for fading channels. *IEEE Transactions on Communications*, 46(5):595–602, May 1998.

[25] 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*, 38(7):70–77, July 2000.

[26] 1xEV: 1x EVolution, IS-856 TIA/EIA standard. *QUALCOMM Inc.*

[27] H. Zhang. Service disciplines for guaranteed performance service in packet-switching networks. *Proceedings of the IEEE*, 83(10):1374–1396, October 1995.

[28] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queueing algorithm. *Journal of Internetworking Research and Experience*, 1:3–26, 1990.

[29] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the single-node case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, 1993.

[30] P. Viswanath, D. Tse, and R. Laroia. Opportunistic beamforming using dumb antennas. *IEEE Transactions on Information Theory*, 48(6):1277–1294, 2002.

[31] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 1996.

[32] G. Stuber. *Principles of Mobile Communication*. Kluwer Academic Publishers, 1996.

[33] G. Hess. *Handbook of Land-mobile Radio System Coverage*. Artech House Publishers, 1998.

[34] M. Gudmundson. Correlation model for shadow fading in mobile radio systems. *Electronics Letters*, 27:2145–2146, November 1991.

[35] E. Gilbert. Capacity of a burst-noise channel. *Bell Systems Technical Journal*, 39:1253–1266, 1960.

[36] S. Lu, V. Bharghavan, and R. Srikant. Fair scheduling in wireless packet networks. *IEEE/ACM Transactions on Networking*, 7(4):473–489, August 1999.

[37] 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.

[38] T. Ng, I. Stoica, and H. Zhang. Packet fair queueing algorithms for wireless networks with location-dependent errors. In *Proceedings of IEEE INFOCOM 1998*, volume 3, New York, 1998.

[39] 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.

[40] Y. Cao and V. Li. Scheduling algorithms in broadband wireless networks. *Proceedings of the IEEE*, 89(1):76–87, January 2001.

[41] 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*, volume 3, 2000.

[42] F. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.

[43] 1x high data rate airlink overview. *QUALCOMM Inc.*

[44] J. Holtzman. CDMA forward link waterfilling power control. In *Proceedings of IEEE Vehicular Technology Conference 2000-Spring*, volume 3. IEEE, 2000.

[45] J. Holtzman. Asymptotic analysis of proportional fair algorithm. In *Proceedings of 12th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, volume 2, 2001.

[46] 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*, 39(2):150–153, February 2001.

[47] 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.

[48] 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.

[49] 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.

[50] 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.

[51] S. Borst and P. Whiting. Dynamic rate control algorithms for CDMA throughput optimization. In *Proceedings of IEEE Infocom 2001*, Alaska, April 2001.

[52] R. Berry. *Power and delay trade-offs in fading channels*. PhD thesis, MIT, 2000.

[53] M. Agarwal and A. Puri. Base station scheduling of requests with fixed deadlines. In *Proceedings of IEEE INFOCOM 2002*. IEEE, 2002.

[54] N. Joshi, S. Kadaba, S. Patel, and G. Sundaram. Downlink scheduling in CDMA data networks. In *Proceedings of ACM Mobicom 2000*, 2000.

[55] D. Zhang and K. Wasserman. Transmission schemes for time-varying wireless channels with partial state observations. In *Proceedings of IEEE INFOCOM 2002*. IEEE, 2002.

[56] S. Oh and K. M. Wasserman. Distributed power control and spreading gain allocation in CDMA data networks. In *1999 IEEE Wireless Communications and Networking Conference*, volume 1. IEEE, 1999.

[57] S. Ramakrishna and J. Holtzman. A scheme for throughput maximization in a dual-class CDMA system. *IEEE Journal on Selected Areas in Communications*, 16(6):830–844, August 1998.

[58] S. Hanly. *Information Capacity of Radio Networks*. PhD thesis, University of Cambridge, 1993.

[59] N. Bambos. Toward power-sensitive network architectures in wireless communications:concepts, issues, and design aspects. *IEEE Personal Communications*, 5(3):50–59, June 1998.

[60] G. Bianchi and A. T. Campbell. A programmable MAC framework for utility-based adaptive quality of service support. *IEEE Journal on Selected Areas in Communications*, 18(2):244–255, February 2000.

[61] C. U. Saraydar, N. B. Mandayam, and D. J. Goodman. Pareto efficiency of pricing-based power control in wireless data networks. In *1999 IEEE Wireless Communications and Networking Conference*, volume 1, 1999.

[62] H. Kushner and G. Yin. *Stochastic Approximation Algorithms and Applications*. Springer-Verlag New York, 1997.

[63] I. Wang, E. Chong, and S. Kulkarni. Weighted averaging and stochastic approximation. *Mathematics of Control, Signals, and Systems*, 1(10):41–60, 1997.

[64] X. Liu, E. Chong, and N. Shroff. Opportunistic transmission scheduling with resource-sharing constraints in wireless networks. *IEEE Journal on Selected Areas in Communications,*, 19(10), October 2001.

[65] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, 1987.

[66] A. Furuskar, M. Frodigh, H. Olofsson, and J. Skold. System performance of EDGE, a proposal for enhanced data rates in existing digital cellular systems. In *Proceedings of 48th IEEE Vehicular Technology Conference*, volume 2. IEEE, 1998.

[67] R. Givan, E. Chong, and H. Chang. Scheduling multiclass packet streams to minimize weighted loss. 2001. manuscript submitted, available on http://www.engr.colostate.edu/∼echong/pubs/cmto.pdf.

[68] H. S. Chang, R. Givan, and E. K. P. Chong. On-line scheduling via sampling. In *Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling (AIPS2000)*, Breckenridge, CO, April 14–17 2000.

[69] J. Peha. Heterogeneous-criteria scheduling: minimizing weighted number of tardy jobs and weighted completion time. *Computers Operations Research*, 22(10):1089–1100, 1995.

[70] J. Peha and F. Tobagi. Evaluating scheduling algorithms for traffic with heterogeneous performance objectives. In *Proceedings of IEEE GLOBECOM*, volume 1, 1990.

[71] B. Hajek and P. Seri. On causal scheduling of multiclass traffic with deadlines. In *Proceedings IEEE Int. Symposium on Information Theory*, 1998.

[72] P. Chow, J. Cioffi, and J. Bingham. A practical discrete multitone transceiver loading algorithm for data transmission over spectrally shaped channels. *IEEE Transactions on Communications*, 43(2/3/4):773–775, 1995.

[73] R. Gallager. *Information Theory and Reliable Communication*. Wiley, 1968.

[74] R. Grunheid, E. Bolinth, and H. Rohling. A blockwise loading algorithm for the adaptive modulation technique in OFDM systems. In *Proceedings of IEEE Vehicular Technology Conference 2001-Fall*, 2001.

[75] S. Lai, R. Cheng, K. Letaief, and C. Tsui. Adaptive tracking of optimal bit and power allocation for OFDM systems in time-varying channels. In *Proceedings of Wireless Communications and Networking Conference*, volume 2. IEEE, 1999.

[76] L. Piazzo. Fast algorithm for power and bit allocation in OFDM systems. *Electronics Letters*, 35(25):2173–2174, December 1999.

[77] Q. Su and S. Schwartz. Effects of imperfect channel information on adaptive loading gain of OFDM. In *Proceedings of IEEE Vehicular Technology Conference 2001-Fall*, 2001.

[78] W. Yu, G. Ginis, and J. Cioffi. An adaptive multiuser power control algorithm for VDSL. In *Proceedings of the IEEE Global Telecommunications Conference*, 2001.

[79] C. Wong, R. Cheng, K. Letaief, and R. Murch. Multiuser OFDM with adaptive subcarrier, bit and power allocation. *IEEE Journal on Selected Areas in Communications*, 17(10):1747–1758, 1999.

[80] H. Yin and H. Liu. An efficient multiuser loading algorithm for OFDM-based broadband wireless systems. In *Proceedings of IEEE Global Telecommunications Conference*, volume 1, 2000.

[81] W. Rhee and J. Cioffi. Increase in capacity of multiuser OFDM system using dynamic subchannel allocation. In *Proceedings of IEEE Vehicular Technology Conference 2000-Spring*, 2000.

[82] C. Yih and E. Geraniotis. Adaptive modulation, power allocation and control for OFDM wireless networks. In *Proceedings of the 11th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, volume 2, 2000.

[83] M. Wahlqvist, H. Olofsson, M. Ericson, C. Ostberg, and R. Larsson. Capacity comparison of an OFDM based multiple access system using different dynamic resource allocation. In *Proceedings of IEEE Vehicular Technology Conference 1997*, volume 3, 1997.

[84] G. Lupo, R. Cautelier, M. Benedetto, and F. Malena. Dynamic resource allocation with a soft handover procedure for application in a broadband system. In *Proceedings of IEEE Vehicular Technology Conference 1999-Fall*, volume 4, 1999.

[85] J. Chuang, N. Sollengerger, and L. Cimini. Power control for dynamic packet assignment in advanced cellular internet service. In *Proceedings of IEEE 48th Vehicular Technology Conference*, volume 3, 1998.

APPENDIX A

A.1 Proof of Prop. 1

Given the fairness requirements $r_1, \cdots, r_N$, the scheduling policy (3.2) is an opportunistic solution of (3.1). In the following, we prove the optimality of our scheduling policy (3.2).

To make the proof easy to understand and provide insight into the scheduling policy, we consider the special case where there are only two users and $P\{U_1 + v^* = U_2\} = 0$. For this special case, the opportunistic scheduling policy is given by $Q^*(\vec{U}) = \text{argmax}(U_1 + v^*, U_2)$, which is illustrated in Figure A.1. Above the line $U_1 + v^* = U_2$, we have $Q^*(\vec{U}) = 2$, while below the line, we have $Q^*(\vec{U}) = 1$. The probability measure of the line is 0 in this case. We show that $E\left(U_{Q^*(\vec{U})}\right) \geq E\left(U_{Q'(\vec{U})}\right)$ for any feasible policy Q' (recall that a feasible policy is a policy that satisfies the time-fraction assignment constraint).

Consider a feasible policy Q' that is different from policy Q^* in regions A and B as shown in Figure A.1, where A, B, C, and D are events given by:

$$
\begin{aligned}
A &= \{Q^*(\vec{U}) = 2, Q'(\vec{U}) = 1\}; \\
C &= \{Q^*(\vec{U}) = Q'(\vec{U}) = 2\}; \\
B &= \{Q^*(\vec{U}) = 1, Q'(\vec{U}) = 2\}; \\
D &= \{Q^*(\vec{U}) = Q'(\vec{U}) = 1\}.
\end{aligned}
$$

It is obvious that $P(A) = P(B)$. If $P(A) = 0$, then $Q'(\vec{U})$ is equal to $Q^*(\vec{U})$ with probability 1 and $E\left(U_{Q'(\vec{U})}\right) = E\left(U_{Q^*(\vec{U})}\right)$ because U_i is bounded. We next show that if $P(A) > 0$, then $E\left(U_{Q^*(\vec{U})}\right) > E\left(U_{Q'(\vec{U})}\right)$. Indeed,

$$
E\left(U_{Q^*(\vec{U})}\right) = \int_D U_1 \, \mathrm{d}P + \int_B U_1 \, \mathrm{d}P + \int_A U_2 \, \mathrm{d}P + \int_C U_2 \, \mathrm{d}P
$$

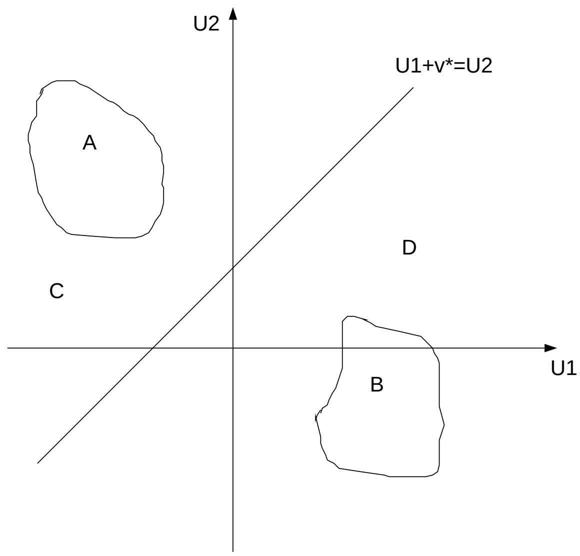

Fig. A.1. Illustration of policies Q^* and Q'; $Q^*(\vec{U}) = 1$ in B and D, and $Q'(\vec{U}) = 1$ in A and D; $Q^*(\vec{U}) = 2$ in A and C, and $Q'(\vec{U}) = 1$ in B and C.

$$E\left(U_{Q'(\vec{U})}\right) = \int_D U_1 \, \mathrm{d}P + \int_B U_2 \, \mathrm{d}P + \int_A U_1 \, \mathrm{d}P + \int_C U_2 \, \mathrm{d}P.$$

Hence,

$$E\left(U_{Q^*(\vec{U})}\right) - E\left(U_{Q'(\vec{U})}\right) = \int_B (U_1 - U_2) \, \mathrm{d}P + \int_A (U_2 - U_1) \, \mathrm{d}P.$$

On A, $U_2 > U_1 + v^*$, so $U_2 - U_1 > v^*$; and on B, $U_2 < U_1 + v^*$, so $U_1 - U_2 > -v^*$. So

$$E\left(U_{Q^*(\vec{U})}\right) - E\left(U_{Q'(\vec{U})}\right) > \int_B -v^* \, \mathrm{d}P + \int_A v^* \, \mathrm{d}P = 0.$$

Hence, $E\left(U_{Q^*(\vec{U})}\right) \geq E\left(U_{Q'(\vec{U})}\right)$ for any feasible policy Q'. The key insight in the above argument is that any policy Q' that schedules user 1 instead of user 2 in any region above the line will lose out to policy Q^* because the performance over that region is inferior for user 1 relative to user 2.

Next, we prove the scheduling policy is optimal *in general*; i.e., for the case of multiple users. Let $Q^*(\vec{U})$ be the opportunistic policy in (3.2). We show that

$E\left(U_{Q^*(\vec{U})}\right) \geq E\left(U_{Q'(\vec{U})}\right)$ for any feasible policy Q'. A similar argument to the one used above applies here.

Let C_{ij} denote the event:

$$C_{ij} = \{Q^*(\vec{U}) = i, Q'(\vec{U}) = j\}$$

for $i = 1, \cdots, N$ and $j = 1, \cdots, N$. Because both Q^* and Q' are feasible policies, we have

$$\begin{aligned} B_j &= \bigcup_{i=1}^{N} C_{ij} = \{Q'(\vec{U}) = j\} \quad \Rightarrow \quad P(B_j) = r_j; \\ A_i &= \bigcup_{j=1}^{N} C_{ij} = \{Q^*(\vec{U}) = i\} \quad \Rightarrow \quad P(A_i) = r_i. \end{aligned}$$

Thus,

$$\begin{aligned} & E\left(U_{Q^*(\vec{U})}\right) - E\left(U_{Q'(\vec{U})}\right) \\ = & \sum_{i=1}^{N}\sum_{j=1}^{N}\int_{C_{ij}} U_{Q^*(\vec{U})}\,\mathrm{d}P - \sum_{i=1}^{N}\sum_{j=1}^{N}\int_{C_{ij}} U_{Q'(\vec{U})}\,\mathrm{d}P \\ = & \sum_{i=1}^{N}\sum_{j=1}^{N}\int_{C_{ij}} U_{Q^*(\vec{U})} - U_{Q'(\vec{U})}\,\mathrm{d}P \\ = & \sum_{i=1}^{N}\sum_{j=1}^{N}\int_{C_{ij}} U_i - U_j\,\mathrm{d}P. \end{aligned}$$

On C_{ij}, we have $Q^*(\vec{U}) = i$ and $Q'(\vec{U}) = j$. Because $Q^*(\vec{U}) = i$, then $U_i + v_i^* \geq U_j + v_j^*$ by the construction of the opportunistic policy Q^*, and so $U_i - U_j \geq v_j^* - v_i^*$. Therefore,

$$\begin{aligned} & E\left(U_{Q^*(\vec{U})}\right) - E\left(U_{Q'(\vec{U})}\right) \\ \geq & \sum_{i=1}^{N}\sum_{j=1}^{N}\int_{C_{ij}} v_j^* - v_i^*\,\mathrm{d}P \\ = & \sum_{j=1}^{N} v_j^* \sum_{i=1}^{N}\int_{C_{ij}} \mathrm{d}P - \sum_{i=1}^{N} v_i^* \sum_{j=1}^{N}\int_{C_{ij}} \mathrm{d}P \\ = & \sum_{j=1}^{N} v_j^* \int_{B_j} \mathrm{d}P - \sum_{i=1}^{N} v_i^* \int_{A_i} \mathrm{d}P \end{aligned}$$

$$
\begin{aligned}
&= \sum_{j=1}^{N} v_j^* r_j - \sum_{i=1}^{N} v_i^* r_i \\
&= 0
\end{aligned}
$$

Hence, $E\left(U_{Q^*(\vec{U})}\right) - E\left(U_{Q'(\vec{U})}\right) \geq 0$ for any feasible policy Q', which completes the proof.

A.2 Improvments of Individual Users

When users have independent performance values, the opportunistic scheduling scheme improves performance for each individual user. Here we prove Prop. 2. Note that Prop. 4 is a trivial extension of Prop. 2.

Denote $T_i = E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right)$, $T_i^r = r_i E(U_i)$, and $\sum_{i=1}^N r_i = 1$. Hence, T_i is the average performance of user i in the opportunistic scheduling policy, and T_i^r is the average performance of a non-opportunistic scheduling policy. Because a non-opportunistic scheduling policy schedules users without considering channel conditions, and r_i is the portion of time slots assigned to user i, we have $T_i^r = r_i E(U_i)$.

In the following, we will show that $T_i \geq T_i^r$ for all i, if U_i is independent of U_j for all $i \neq j$.

$$
\begin{aligned}
T_i &= E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right) \\
&= E(U_i | Q(\vec{U}) = i) P(Q(\vec{U}) = i) \\
&= r_i E(U_i | Q(\vec{U}) = i) \\
T_i^r &= r_i E(U_i)
\end{aligned}
$$

Hence, to prove that $T_i \geq T_i^r$, we only need to prove that

$$E(U_i | Q(\vec{U}) = i) \geq E(U_i)$$

w.p.1 (with probability 1). We have

$$
\begin{aligned}
Q(\vec{U}) = i \quad &\Rightarrow \quad U_i + v_i^* \geq U_j + v_j^*, \quad \text{for all } j \\
&\Leftrightarrow \quad U_i \geq \max_{j \neq i}(U_j + v_j^*) - v_i^*
\end{aligned}
$$

Let $Y = \max_{j \neq i}(U_j + v_j^*) - v_i^*$, we have

$$\begin{aligned} E(U_i|Q(\vec{U}) = i) &\geq E(U_i|U_i \geq Y) \\ &= E(U_i|U_i \geq Y)P(U_i \geq Y) + E(U_i|U_i \geq Y)P(U_i < Y) \\ &\geq E(U_i|U_i \geq Y)P(U_i \geq Y) + E(U_i|U_i < Y)P(U_i < Y) \\ &= E(U_i); \end{aligned}$$

i.e., $E(U_i|Q(\vec{U}) = i) \geq E(U_i)$. Hence, $T_i \geq T_i^r$. $\square$

Note that in general we cannot prove that $T_i \geq T_i^r$. Here is a counter example where two users have a negative correlation coefficient.

$$\begin{aligned} U_1 &= \begin{cases} 100 & \text{with prob. } 0.1 \\ 3 & \text{with prob. } 0.5 \\ 1 & \text{with prob. } 0.4 \end{cases} \\ U_2 &= \begin{cases} -1 & \text{if } \;\; U_1 = 100 \\ 2 & \text{if } \;\; U_1 = 3 \\ 1 & \text{if } \;\; U_1 = 1 \end{cases} \end{aligned}$$

Furthermore, $r_1 = r_2 = 0.5$ and

$$\begin{aligned} E(U_1) &= 100 * 0.1 + 3 * 0.5 + 1 * 0.4 = 11.9 \\ E(U_2) &= -1 * 0.1 + 2 * 0.5 + 1 * 0.4 = 1.3 \\ E(U_1^2) &= 1004.9 \\ E(U_2^2) &= 2.5 \\ \sigma_1 &= \left(E(U_1^2) - E^2(U_1)\right)^{1/2} = 29.3 \\ \sigma_2 &= \left(E(U_2^2) - E^2(U_2)\right)^{1.2} = 0.9 \\ \rho &= \frac{E(U_1U_2) - E(U_1)E(U_2)}{\sigma_1\sigma_2} = -0.84. \end{aligned}$$

It is obvious that $\vec{v}^* = [1, 0]$ and when $U_1 = 3$ and $U_2 = 2$, we let user 1 transmit with probability 0.8. Hence,

$$T_2 = E\left(U_2\mathbf{1}_{\{Q=2\}}\right) = 2 * 0.5 * (1 - 0.8) + 1 * 0.4 = 0.6 < E(U_2) * r_2 = 0.65.$$

One intuition from this counter example is: correlation coefficient is not an accurate measure of users' dependence on each other. In this case, negative correlation coefficient does not imply that when one user is good, the other one is bad.

A.3 Proof of Prop. 5

Recall that the performance-based scheduling problem is defined as

$$\begin{array}{lll} \underset{Q\in\Theta}{\text{maximize}} & E\left(U_{Q(\vec{U})}\right) & \\ \text{subject to} & E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right) \geq a_i E\left(U_{Q(\vec{U})}\right), & i = 1, 2, \cdots, N. \end{array}$$

The policy Q^* is defined as

$$Q^*(\vec{U}) = \underset{i}{\text{argmax}}\left((\kappa + \nu_i^*)\,U_i\right),$$

where $\kappa = 1 - \sum_{i=1}^N a_i \nu_i^*$, and the ν_i^*'s are chosen so that:

1. $\min_i(\nu_i^*) = 0$
2. $E\left(U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) \geq a_i E\left(U_{Q^*(\vec{U})}\right)$ for all i
3. For all i, if $E\left(U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) > a_i E\left(U_{Q^*(\vec{U})}\right)$, then $\nu_i^* = 0$.

In the following, we prove that Q^* is an optimal policy.

Proof: Let Q be a policy satisfying $E(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}) \geq a_i E(U_{Q(\vec{U})})$ for all i. Because $\nu_i^* \geq 0$,

$$\begin{aligned} E\left(U_{Q(\vec{U})}\right) &\leq E\left(U_{Q(\vec{U})}\right) + \sum_{i=1}^N \nu_i^* \left(E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right) - a_i E\left(U_{Q(\vec{U})}\right)\right) \\ &= \left(1 - \sum_{i=1}^N \nu_i^* a_i\right) E\left(\sum_{i=1}^N U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right) + \sum_{i=1}^N \nu_i^* E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right) \\ &= \sum_{i=1}^N (\kappa + \nu_i^*) E\left(U_i \mathbf{1}_{\{Q(\vec{U})=i\}}\right). \end{aligned}$$

By the definition of Q^*, we have

$$\sum_{i=1}^N \left((\kappa + \nu_i^*) U_i\right) \mathbf{1}_{\{Q(\vec{U})=i\}} \leq \sum_{i=1}^N \left((\kappa + \nu_i^*) U_i\right) \mathbf{1}_{\{Q^*(\vec{U})=i\}}$$

Hence,

$$
\begin{aligned}
E\left(U_{Q(\vec{U})}\right) &\leq \sum_{i=1}^{N}(\kappa+\nu_i^*)E\left(U_i\mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) \\
&= E\left(U_{Q^*(\vec{U})}\right)+\sum_{i=1}^{N}\nu_i^*\left(E\left(U_i\mathbf{1}_{\{Q^*(\vec{U})=i\}}\right)-a_iE\left(U_{Q^*(\vec{U})}\right)\right) \\
&= E\left(U_{Q^*(\vec{U})}\right),
\end{aligned}
$$

which completes the proof. □

A.4 Proof of Prop. 6

The minimum-performance guarantee problem is formulated as:

$$
\begin{aligned}
&\underset{Q}{\text{maximize}} && E\left(U_{Q(\vec{U})}\right) \\
&\text{subject to} && E\left(U_i\mathbf{1}_{\{Q(\vec{U})=i\}}\right) \geq C_i, \qquad i=1,2,\cdots,N,
\end{aligned}
$$

where $E(U_i\mathbf{1}_{\{Q(\vec{U})=i\}})$ is the average performance of user i, and $C_i \geq 0$ is the minimum performance requirement of user i. The policy Q^* is defined as:

$$
Q^*(\vec{U}) = \underset{i}{\operatorname{argmax}}(\alpha_i^* U_i), \tag{1}
$$

where the α_i^*'s are chosen so that:

1. $\min_i(\alpha_i^*)=1$
2. $E\left(U_i\mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) \geq C_i$ for all i
3. For any user i, if $E\left(U_i\mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) > C_i$, then $\alpha_i^*=1$.

In the following, we prove that Q^* is an optimal policy.

Proof: Let Q be a policy satisfying $E(U_i\mathbf{1}_{\{Q(\vec{U})=i\}}) \geq C_i$ for all i. Because $\alpha_i^* \geq 1$,

$$
\begin{aligned}
E\left(U_{Q(\vec{U})}\right) &\leq E\left(U_{Q(\vec{U})}\right)+\sum_{i=1}^{N}(\alpha_i^*-1)\left(E\left(U_i\mathbf{1}_{\{Q(\vec{U})=i\}}\right)-C_i\right) \\
&= \sum_{i=1}^{N}E\left(U_i\mathbf{1}_{\{Q(\vec{U})=i\}}\right)+\sum_{i=1}^{N}(\alpha_i^*-1)E\left(U_i\mathbf{1}_{\{Q(\vec{U})=i\}}\right)-\sum_{i=1}^{N}(\alpha_i^*-1)C_i \\
&= \sum_{i=1}^{N}E\left(\alpha_i^*U_i\mathbf{1}_{\{Q(\vec{U})=i\}}\right)-\sum_{i=1}^{N}(\alpha_i^*-1)C_i.
\end{aligned}
$$

By the definition of Q^*, we have

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

Hence,

$$\begin{aligned} E\left(U_{Q(\vec{U})}\right) &\leq \sum_{i=1}^{N} E\left(\alpha_i^* U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) - \sum_{i=1}^{N} (\alpha_i^* - 1) C_i \\ &= E\left(U_{Q^*(\vec{U})}\right) + \sum_{i=1}^{N} (\alpha_i^* - 1)\left(E\left(U_i \mathbf{1}_{\{Q^*(\vec{U})=i\}}\right) - C_i\right) \\ &= E\left(U_{Q^*(\vec{U})}\right), \end{aligned}$$

which completes the proof. □

A.5 Proof of Existence

In this section, we prove that there exists an optimal policy as defined in (4.2) for the temporal fairness scheduling problem. In other words, there exist v_i^*s that satisfy the following conditions:

1. $\min_i(v_i^*) = 0$
2. $P\{Q^*(\vec{U}) = i\} \geq r_i$ for all i
3. For all i, if $P\{Q^*(\vec{U}) = i\} > r_i,$ then $v_i^* = 0,$

where Q^* is defined as

$$Q^*(\vec{U}) = \underset{i}{\operatorname{argmax}}(U_i + v_i^*).$$

We first prove the existence of such a policy for the special case studied in Chapter 2, where $\sum_{i=1}^{N} r_i = 1$. We then generalize the result to the case where $\sum_{i=1}^{N} r_i \leq 1$.

Lemma 1 *For any vector $\vec{r} = \{r_1, r_2, \cdots, r_N\}$, where $\sum_{i=1}^{N} r_i = 1$ and $r_i \geq 0$ for all i, there exists $\vec{v}^*$ and corresponding tie-breaking probability such that $P\{Q^*(\vec{U}) = i\} = r_i$ for all i, where the policy Q^* is defined as:*

$$Q^*(\vec{U}) = \underset{i=1,\cdots,N}{\operatorname{argmax}} U_i + v_i^*,$$

and when a tie occurs, we choose one of the maximizing users with a certain probability.

Proof: 1. Special case: We first consider the case where the U_is are continuous random variables and have bounded joint density functions. We use induction. When there are only two users in the system, it is easy to see that there exists v^*, such that $P\{U_1 + v^* \geq U_2\} = r_1$; i.e, $v_1^* = v^*$ and $v_2^* = 0$.

Suppose the lemma is true when there are $N-1$ users. We will prove it is true for the N-user case. In the N-user case, if $r_i = 0$ for any i, then it degenerates to the $N-1$ user case, which holds by the induction hypothesis. Hence, in the following, we only consider the case where $r_i > 0$ for all i.

Let

$$\tilde{r_i} = r_i + r_N/(N-1), \quad i = 1, \cdots, N-1.$$

By the induction hypothesis, there exists $v_1^0, \cdots, v_{N-1}^0$, such that

$$P\{U_i + v_i^0 \geq \max_{j=1,\cdots,N-1} U_j + v_j^0\} = \tilde{r_i}, \quad i = 1, \cdots, N-1.$$

Because $E(U_i) < \infty$, there exists M such that

$$P\{U_i > M\} < r_i/(N-1), \quad i = 1, \cdots, N.$$

Define

$$f_i(\vec{v}) = P\{U_i + v_i \geq \max_{j \neq i} U_j + v_j\}.$$

Let

$$v_i^1 = v_i^0 - \min_{j=1,\cdots,N-1} v_j^0 + M, \quad i = 1, \cdots, N-1,$$

and $v_N^1 = 0$. Recall that U_is are non-negative. Then we have

$$f_i(\vec{v}^1) \geq \tilde{r_i} - \frac{r_N}{N-1} = r_i, \quad i = 1, \cdots, N-1.$$

In the following, we will construct a sequence of $\vec{v}^i$, $i = 1, 2, \cdots$, that converges. For the ith step, we do the following: we pick a user j, where

$$j = (i \bmod (N-1)) + 1.$$

Let $v_k^i = v_k^{i-1}$ for all $k \neq j$, and set the value of v_j^i such that

$$f_j(\vec{v}^i) = r_j.$$

Apparently, $v_j^i \leq v_j^{i-1}$. Because $f_k(\vec{v})$ $(k \neq j)$ increases when v_j decreases, we have

$$f_k(\vec{v}^i) \geq f_k(\vec{v}^{i-1}) \geq r_k, \quad k \neq j.$$

Thus, we have constructed a decreasing sequence $\vec{v}^i$ that has a lower bound $(v_j^i \geq -M)$. Hence, it converges. Let

$$v_j^* = \lim_{i \to \infty} v_j^i, \quad j = 1, \cdots, N-1.$$

and $v_N^* = 0$.

Because the U_js are continuous random variables with bounded joint density function, $f_j(\vec{v})$ is a continuous bounded function of $\vec{v}$. Hence, $f_j(\vec{v}^i)$ converges. Furthermore, there exists a subsequence $f_j(\vec{v}^i) = r_j$, where $i = m(N-1) + j - 1$. Hence,

$$f_j(\vec{v}^*) = r_j, \quad j = 1, \cdots, N-1,$$

and thus,

$$f_N(\vec{v}^*) = 1 - \sum_{j=1}^{N-1} r_j = r_N.$$

In other words, if the U_js are continuous random variable, then there exists $\vec{v}^*$ such that

$$P\{U_i + v_i^* \geq \max_{j \neq i} U_j + v_j^*\} = r_i, \quad i = 1, \cdots, N.$$

2. General Case: Let us consider the N-user case, where $r_i > 0$ for $i = 1, \cdots, N$. To prove Lemma 1, we construct a sequence of continuous random variables with bounded joint density functions to approximate random viables U_is, and obtain v_i^*s by the limit of the parameters of the sequence of the constructed continuous random variables. Then we find tie-breaking probabilities, which concludes the proof.

First, we use a sequence of continuous random variables $\{U_i^n\}$ to approximate random variables U_is:

$$U_i^n = U_i + \varepsilon_i^n, \quad j = 1, \cdots, N,$$

where ε_i^n is a random variable uniformly distributed in $[0, 1/n]$, ε_i^ns are independent, and ε_i^n is independent of U_i. Next, we show that U_i^n have bounded joint density function.

$$\begin{aligned}
&P\{x_1 - dx_1 \le U_1^n \le x_1, \cdots, x_N - dx_N \le U_N^n \le x_N\} \\
= &\int_{x_1-dx_1-1/n}^{x_1-dx_1} \cdots \int_{x_N-dx_N-1/n}^{x_N-dx_N} dF(U_1, \cdots, U_N) \int_{x_1-dx_1-U_1}^{x_1-U_1} nd\varepsilon_1^n \cdots \int_{x_N-dx_N-U_N}^{x_N-U_N} nd\varepsilon_N^n \\
= &n^N dx_1 \ldots dx_N \int_{x_1-dx_1-1/n}^{x_1-dx_1} \cdots \int_{x_N-dx_N-1/n}^{x_N-dx_N} dF(U_1, \cdots, U_N)
\end{aligned}$$

Thus, the joint density function of U_i^n satisfies

$$\begin{aligned}
p_{U_1^n, \ldots U_N^n}(x_1, \ldots x_n) &= n^N \int_{x_1-dx_1-1/n}^{x_1-dx_1} \cdots \int_{x_N-dx_N-1/n}^{x_N-dx_N} dF(U_1, \cdots, U_N) \\
&\le n^N.
\end{aligned}$$

Hence, U_i^ns have bounded joint density functions. Let $\vec{v}^{n*}$ satisfy

$$P\{U_i^n + v_i^{n*} \ge \max_{j \ne i}(U_j^n + v_j^{n*})\} = r_i, \quad i = 1, \cdots, N,$$

where $v_N^{n*} = 0$. Note that the sequence $\{\vec{v}^{n*}\}$ is bounded. Thus, there exists a subsequence that converges; i.e.,

$$\lim_{k \to \infty} v_i^{n_k*} = v_i^*, \quad i = 1, \cdots, N.$$

In the following, we focus only on this convergent subsequence.

For notational convenience, we write

$$\begin{aligned}
W_i &= U_i + v_i^*, \quad i = 1, \cdots, N, \\
W_i^k &= U_i^{n_k} + v_i^{n_k*}, \quad i = 1, \cdots, N, \\
r_i' &= P\{W_i > \max_{j \ne i} W_j\}, \quad i = 1, \cdots, N, \\
\Delta r_i &= r_i - r_i', \quad i = 1, \cdots, N.
\end{aligned}$$

Let Y be a random varialbe. Then, $\lim_{\epsilon \to 0} P\{0 < Y < \epsilon\} = 0$. Hence, for any $\delta > 0$, there exists $\epsilon > 0$ such that $P\{0 < Y < 10\epsilon\} < \delta$. In other words, for any $\delta > 0$, there exists $\epsilon > 0$ such that

$$P\{\max_{j \ne i} W_j < W_i < \max_{j \ne i} W_j + 10\epsilon\} < \delta, \quad i = 1, \cdots, N. \tag{2}$$

$$P\{W_i < \max_{j\neq i} W_j < W_i + 10\epsilon\} < \delta, \quad i = 1, \cdots, N. \tag{3}$$

For this ϵ, there exists K such that for any $k > K$, we have $1/n_k < \epsilon$ and

$$|v_i^{n_k *} - v_i^*| < \epsilon, \quad i = 1, \cdots, N.$$

Next, we will show that

$$r_i' - \delta \le P\{W_i^k \ge \max_{j\neq i} W_j^k + 4\epsilon\} \le r_i'.$$

If $W_i^k \ge \max_{j\neq i} W_j^k + 4\epsilon$, then we have

$$\begin{aligned} & U_i + \frac{1}{n_k} + v_i^* + \epsilon \ge W_i^k \ge \max_{j\neq i} W_j^k + 4\epsilon \ge \max_{j\neq i}(U_j + v_j^* - \epsilon) + 4\epsilon \\ \Rightarrow \quad & U_i + v_i^* \ge \max_{j\neq i} U_j + v_j^* + \epsilon > \max_{j\neq i} U_j + v_j^* \\ \Rightarrow \quad & W_i > \max_{j\neq i} W_j. \end{aligned}$$

Note that $P(A) \le P(B)$ if the occurrence of event A guarantees the occurrence of event B. Using this fact, we have

$$P\{W_i^k \ge \max_{j\neq i} W_j^k + 4\epsilon\} \le P\{W_i > \max_{j\neq i} W_j\} = r_i'.$$

Furthermore, if $W_i > \max_{j\neq i} W_j + 8\epsilon$, then we have

$$\begin{aligned} & U_i^{n_k} + v_i^{n_k *} + \epsilon \ge U_i + v_i^* \ge \max_{j\neq i} U_j^{n_k} - \epsilon + v_j^{n_k *} - \epsilon + 8\epsilon \\ \Rightarrow \quad & W_i^k \ge \max_{j\neq i} W_j^k + 4\epsilon. \end{aligned}$$

Hence,

$$\begin{aligned} P\{W_i^k \ge \max_{j\neq i} W_j^k + 4\epsilon\} &\ge P\{W_i > \max_{j\neq i} W_j + 8\epsilon\} \\ &= P\{W_i > \max_{j\neq i} W_j\} - P\{\max_{j\neq i} W_j < W_i \le \max_{j\neq i} W_j + 8\epsilon\} \\ &\ge r_i' - \delta. \end{aligned}$$

The second inequality is due to (2). In summary, we have

$$r_i' - \delta \le P\{W_i^k \ge \max_{j\neq i} W_j^k + 4\epsilon\} \le r_i', \quad i = 1, \cdots, N. \tag{4}$$

In the following, we are going to construct tie-breaking probabilities. Let M_i^k be the event $\{\max_{j\neq i} W_j^k \leq W_i^k < \max_{j\neq i} W_j^k + 4\epsilon\}$. We have

$$r_i = P\{W_i^k \geq \max W_j^k\} = P\{M_i^k\} + P\{W_i^k \geq \max_{j\neq i} W_j^k + 4\epsilon\}, \quad i = 1, \cdots, N.$$

Substitude it into (4), we have

$$\Delta r_i \leq P\{M_i^k\} \leq \Delta r_i + \delta. \tag{5}$$

Let x be an event that tie occurs and C_x be the set of users that achieve the maximum value of $U_i + v_i^*$; i.e., $C_x = \{i_1, i_2, \cdots, i_l\}$, where $l \geq 2$, and x is the event

$$U_{i_1} + v_{i_1}^* = U_{i_2} + v_{i_2}^* = \cdots = U_{i_l} + v_{i_l}^* > \max_{i \notin C_x} U_i + v_i^*.$$

Let $X_i = \{x : x \in X \cap i \in C_x\}$; i.e., X_i is the set of tie events where user i is one of the users that achieve the maximum. We have

$$P\{M_i^k\} = \sum_{x \in X_i} P\{M_i^k | x\} P\{x\} + P\{M_i^k \cap W_i > \max_{j\neq i} W_j\} + P\{M_i^k \cap W_i < \max_{j\neq i} W_j\}. \tag{6}$$

Note that

$$\begin{aligned}
& P\{M_i^k \cap W_i > \max_{j\neq i} W_j\} \\
= \;& P\{M_i^k \cap W_i > \max_{j\neq i} W_j + 8\epsilon\} + P\{M_i^k \cap \max_{j\neq i} W_j < W_i \leq \max_{j\neq i} W_j + 8\epsilon\} \\
\leq \;& P\{M_i^k \cap W_i^k > \max_{j\neq i} W_j^k + 4\epsilon\} + P\{\max_{j\neq i} W_j < W_i \leq \max_{j\neq i} W_j + 8\epsilon\} \\
\leq \;& P\{\max_{j\neq i} W_j < W_i \leq \max_{j\neq i} W_j + 8\epsilon\} \\
\leq \;& \delta,
\end{aligned}$$

and

$$\begin{aligned}
& P\{M_i^k \cap W_i < \max_{j\neq i} W_j\} \\
= \;& P\{M_i^k \cap W_i + 8\epsilon < \max_{j\neq i} W_j\} + P\{M_i^k \cap W_i < \max_{j\neq i} W_j \leq W_i + 8\epsilon\} \\
\leq \;& P\{M_i^k \cap W_i^k + 4\epsilon < \max_{j\neq i} W_j^k\} + P\{W_i < \max_{j\neq i} W_j \leq W_i + 8\epsilon\} \\
\leq \;& P\{W_i < \max_{j\neq i} W_j \leq W_i + 8\epsilon\} \\
\leq \;& \delta.
\end{aligned}$$

Substitute these two inequalities into (6) to get

$$P\{M_i^k\} - 2\delta \le \sum_{x \in X} P\{M_i^k|x)P(x) \le P\{M_i^k\}.$$

Substituting the above into (5),

$$\Delta r_i - 2\delta \le \sum_{x \in X} P\{M_i^k|x\}P(x) \le \Delta r_i + \delta.$$

Because we chose δ arbitrarily, we have

$$\lim_{k\to\infty} \sum_{x \in X} P\{M_i^k|x\}P(x) = \Delta r_i.$$

Because $P\{M_i^k|x\}$ is bounded, we have a subsequence such that

$$\lim_{l\to\infty} P\{M_i^{k_l}|x\} = p_i(x), \quad i \in C_x,$$

and $p_i(x) = 0$ if $i \notin C_x$.

Because

$$\sum_x p_i(x)P(x) = \Delta r_i.$$

we have

$$P\{U_i + v_i^* > \max_{j\neq i} U_j + v_j^*\} + \sum_x p_i(x)P(x) = \Delta r_i + r_i' = r_i.$$

Hence, there exist $\vec{v}^*$ and a tie-breaking probability such that

$$P\{Q^*(\vec{U})\} = r_i, \quad i = 1, \cdots, N,$$

where

$$Q^*(\vec{U}) = \operatorname*{argmax}_{i=1,\cdots,N} U_i + v_i^*,$$

and when a tie occurs (event x), the policy Q^* chooses user i with probability $p_i(x)$. $\square$

So far, we have proved Lemma 1. Recall that in the resource-based fairness scheduling problem, we define the policy Q^* as follows:

$$Q^*(\vec{U}) = \operatorname*{argmax}_i (U_i + v_i^*), \tag{7}$$

where the v_i^*s are chosen such that:

1. $\min_i(v_i^*) = 0$,

2. $P\{Q^*(\vec{U}) = i\} \geq r_i$ for all i,

3. For all i, if $P\{Q^*(\vec{U}) = i\} > r_i$, then $v_i^* = 0$.

In the following, we will construct v_i^*s that satisfy the above conditions. First, set $\vec{v}^0 = 0$. Let

$$Q(\vec{U}) = \operatorname{argmax} U_i + v_i^0,$$

where ties are broken randomly. If

$$P\{Q(\vec{U}) = i\} \geq r_i, \quad i = 1, \cdots, N,$$

then $\vec{v}^* = \vec{v}^0 = 0$ satisfies all the conditions required. Otherwise, there exists a nonempty set of users, B, such that

$$P\{Q(\vec{U}) = i\} < r_i, \quad i \in B.$$

We can consider B as the set of unfortunate users, which have to take advantage of other users to satisfy their QoS requirements. Let L be the number of users in B. Without loss of generality, we assume that $B = \{1, 2 \cdots, L\}$.

Let

$$\begin{aligned} U'_{L+1} &= \max_{i \notin B} U_i, \\ r'_{L+1} &= 1 - \sum_{i \in B} r_i, \\ U'_i &= U_i, \quad i = 1, \cdots, L, \\ r'_i &= r_i, \quad i = 1, \cdots, L. \end{aligned}$$

Following the lemma, for random variables $U'_1, \cdots, U'_{L+1}$, there exists $\vec{v}'$ and tie-breaking probabilities such that

$$P\{Q(\vec{U}') = i\} = r'_i, \quad i = 1, \cdots, L+1.$$

Note that $v'_{L+1} = \min_{i=1,\cdots,L+1} v'_i$. Otherwise, there exists i such that $P\{Q(\vec{U}') = i\} < r'_i$ for some $i \in B$.

Let the policy Q be

$$Q(\vec{U}) = \begin{cases} Q(\vec{U}') & \text{if } Q(\vec{U}') \leq L, \\ \text{argmax}_{i>L}\, U_i & \text{if } Q(\vec{U}') = L+1, \end{cases}$$

where ties are broken randomly if $Q(\vec{U}') = L+1$. We stop if the following conditions are satisfied:

$$P\{Q(\vec{U}) = i\} = r_i, \quad i \in B$$
$$P\{Q(\vec{U}) = i\} \geq r_i, \quad i \notin B.$$

Then, we set

$$v_j^* = 0, \quad i \notin B$$
$$v_j^* = v_j' - v_{L+1}' \geq 0, \quad i \in B.$$

On the other hand, if there exists users such that

$$P\{Q(\vec{U}) = i\} < r_i,$$

where $i > L$, then we move these users to set B and implement the same procedure. Note that we can have at most N steps because at each iteration the number of users in B increases at least by one. $\square$

A.6 Proof of Convexity

We next prove Prop. 7. Suppose $\vec{C}^1$ and $\vec{C}^2$ are two achievable performance vectors. We will show that for any $0 \leq a \leq 1$, $\vec{C} = a\vec{C}^1 + (1-a)\vec{C}^2$ is also achievable. Let $\vec{C}^1$ be achieved by policy Q^1 and $\vec{C}^2$ be achieved by policy Q^2. Define a "combined" policy Q by

$$Q(\vec{U}) = \begin{cases} Q^1(\vec{U}) & \text{if } A = 1 \\ Q^2(\vec{U}) & \text{otherwise,} \end{cases}$$

where A is a random variable such that $P\{A = 1\} = a$. Then $\vec{C}$ is achieved by policy Q. $\square$

A.7 Proof of Asymptotic Performance Bound

In this section, we prove the asymptotic results presented in Prop. 8.

Proof: It is straight-forward to show that $E(Z_n) = O(n)$. We have

$$\begin{aligned} Z_n &= \max_i U_i \leq |U_1| + \cdots + |U_N| \\ \Rightarrow \quad E(Z_n) &\leq E(|U_1|) + \cdots + E(|U_n|) \leq Cn. \end{aligned}$$

Hence, $E(Z_n) = O(n)$.

Next, we prove the second part of the proposition. For any $\epsilon_0 > 0$, there exists α such that $1 < \alpha < 1/(1-\epsilon_0)$. Let

$$F(x) = \begin{cases} 1 - \frac{1}{x^\alpha} & x \geq 1 \\ 0 & x < 1. \end{cases}$$

We will show that if the U_is are i.i.d. random variables with the above distribution function F, then

$$\lim_{n\to\infty} \frac{E(Z_n)}{n^{1-\epsilon}} = \infty.$$

Let

$$\begin{aligned} b_n &= n^{1/\alpha}, \\ Y_n &= Z_n/b_n, \\ H(x) &= \begin{cases} \exp(-x^{-\alpha}) & \text{if } x > 0, \\ 0 & \text{otherwise}, \end{cases} \\ E(W) &= \int_0^\infty 1 - H(x)\,\mathrm{d}x, \end{aligned}$$

where W is a random variable with the distribution function $H(x)$. Note that $0 < E(W) < \infty$. Let

$$\begin{aligned} F_n(x) &= P\{Y_n \leq x\} \\ &= P\{Z_n \leq b_n x\} \\ &= (P\{U_i \leq b_n x\})^n \\ &= \begin{cases} \left(1 - \frac{1}{x^\alpha n}\right)^n & \text{if } x \geq n^{-1/\alpha} \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Because

$$\lim_{n\to\infty}\left(1-\frac{1}{x^{\alpha}n}\right)^{n}=\exp(-x^{-\alpha}),\quad x>0,$$

Y_n converges to W in distribution[1]. Next, we show that $E(Y_n)$ converges to $E(W)$. Note that $\left(1-\frac{1}{x}\right)^{x}$ is an increasing function of x for $x>1$.[2] Hence, for $x>1$ and $n\geq 1$, we have

$$\begin{aligned} F_n(x) &= \left[\left(1-\frac{1}{x^{\alpha}n}\right)^{x^{\alpha}n}\right]^{\frac{1}{x^{\alpha}}} \\ &\geq \left(1-\frac{1}{x^{\alpha}}\right)^{x^{\alpha}\frac{1}{x^{\alpha}}} \\ &= F(x). \end{aligned}$$

Let $\epsilon>0$ be given. Because U_i has finite mean, there exists $M'>1$ such that

$$\int_{M'}^{\infty}1-F_n(x)\,\mathrm{d}x\leq\int_{M'}^{\infty}1-F(x)\,\mathrm{d}x\leq\epsilon/3. \tag{8}$$

Furthermore, because W has finite mean, there exists $M>M'$ such that

$$\int_{M}^{\infty}1-H(x)\,\mathrm{d}x\leq\epsilon/3.$$

Define function f by

$$f(x)=\begin{cases} M & \text{if } x\geq M,\\ x & \text{otherwise.}\end{cases}$$

Then $f(x)$ is a bounded continuous function. Because Y_n converges to W in distribution,

$$\lim_{n\to\infty}E(f(Y_n))=E(f(W)).$$

Hence, there exists L such that for $n>L$,

$$|E(f(Y_n))-E(f(W))|\leq\epsilon/3.$$

[1]This convergence is also a result due to theories in extreme order statistics. We prove it here for completeness and convenience of readers.

[2]Note that $\ln(1+y)\leq y$ for $y\geq 0$. Let $f(x)=x\ln(1-1/x)$ for $x>1$. Because $f'(x)=\ln(1-1/x)+1/(x-1)=1/(x-1)-\ln(1+1/(x-1))\geq 0$, $f(x)$ is an increasing function of x for $x>1$.

Furthermore, because

$$P(Y_n \leq x) = P(f(Y_n) \leq x), \quad 0 \leq x < M,$$

we have

$$E(Y_n) = E(f(Y_n)) + \int_M^\infty 1 - F_n(x)\,\mathrm{d}x.$$

From (8) and $M' < M$, we have

$$|E(Y_n) - E(f(Y_n))| \leq \epsilon/3.$$

Similarly,

$$|E(f(W)) - E(W)| \leq \epsilon/3.$$

Hence, for the given $\epsilon > 0$, we have that for $n > L$,

$$\begin{aligned} & |E(Y_n) - E(W)| \\ \leq\ & |E(f(Y_n)) - E(f(W))| + |E(f(Y_n)) - E(Y_n)| + |E(f(W)) - E(W)| \\ \leq\ & \epsilon. \end{aligned}$$

Thus,

$$\lim_{n\to\infty} E(Y_n) = E(W);$$

i.e.,

$$\lim_{n\to\infty} \frac{E(Z_n)}{n^{1/\alpha}} = E(W).$$

Recall that $\alpha < 1/(1-\epsilon_0)$. Hence,

$$\lim_{n\to\infty} \frac{E(Z_n)}{n^{1-\epsilon_0}} = \infty.$$

□

VITA

Xin Liu received her B.S. and M.S. degrees in electrical engineering from Xi'an Jiaotong Univerisity, People's Republic of China, in 1994 and 1997, respectively. Since then she has been working toward a Ph.D. degree in the School of Electrical and Computer Engineering at Purdue University. She is a research assistant in the Network Engineering and Wireless Systems group. Her research interest include wireless networks, resource allocation and scheduling, and QoS provisioning.