1

Opportunistic Communication: A System View

Pramod Viswanath

*Department of Electrical and Computer Engineering
University of Illinois, Urbana-Champaign*

The wireless medium is often called a *fading* channel: the pejorative adjective suggests that the intrinsic temporal and frequency variations are an impediment to reliable communication. While not untrue, the channel fluctuations are turned from *foe to friend* in some scenarios. A concrete situation is when the time scale of communication is much larger than that of the channel fluctuations: the so-called ergodic fading channel. In a point-to-point ergodic fading channel, the transmitter can make good use of the channel state information (CSI): by devoting more power to when the channel is good and less (or even none) when the channel is bad, the rate of reliable communication is improved. The improvement is significant when the operating signal to noise ratio (SNR) is small; this is a simple instance of opportunistic communication. Another important instance is in cellular communication: by scheduling user transmissions when their channel conditions are good, the system throughput is improved. This effect is called *multiuser diversity* and its role in cellular communication is the focus of this chapter.

1.1 Multiuser Diversity Gain

Consider the single antenna downlink flat fading channel with K users:

$$y_k[m] = h_k[m]x[m] + w_k[m], \quad k = 1, \dots, K, \quad (1.1)$$

where $\{h_k[m]\}_m$ is the channel fading process of user k . There is an average power constraint of P on the transmit signal and $w_k[m] \sim \mathcal{CN}(0, N_0)$ are i.i.d. in time m (for each user $k = 1, \dots, K$). For concreteness, consider the symmetric case: $\{h_k[m]\}_m$ are identically distributed processes for $k = 1 \dots K$. Further, let us also suppose that the processes $\{h_k[m]\}_m$ are ergodic (i.e., the time average of every realization equals the statistical average).

When the single transmitter tracks the channel fluctuations of all the users (with each user tracking its own channel fluctuation), the sum capacity of the downlink is achieved by transmitting only to the best user; as the channels vary, we can pick the best user at each time and further allocate it an appropriate waterfilling power subject to the average power constraint. The corresponding sum rate can be interpreted as the waterfilling capacity of a point-to-point link with a power constraint equal to the transmit power, and a fading process whose magnitude varies as $\{\max_k |h_k[m]|\}$. Compared to a system with a single transmitting user, the multiuser gain comes from the fact that the effective channel gain at time m that is improved from $|h_1[m]|^2$ to $\max_{1 \leq k \leq K} |h_k[m]|^2$. This effect is entirely due to the ability to dynamically schedule resources among the users as a function of the channel state.

The corresponding full CSI (both at the transmitter and the receivers) sum capacity is increased due to the *multiuser diversity* effect: when there are many users which fade *independently*, at any one time there is a high probability that one of the users will have a strong channel. By allowing only that user to transmit, the shared channel resource is used in the most efficient manner and the total system throughput is maximized. The larger the number of users, the stronger tends to be the strongest channel, and the more the multiuser diversity gain.

The amount of multiuser diversity gain depends crucially on the *tail* of the fading distribution $|h_k|^2$: the heavier the tail, the more likely there is a user with a very strong channel, and the larger the multiuser diversity gain. This is shown in Figure 1.1, where the sum capacity is plotted as a function of the number of users for both Rayleigh and Ricean fading with κ -factor equaling to 5, with the SNR fixed at 0 dB. Ricean fading models the situation when there is a strong specular line-of-sight path plus many small reflected paths. The parameter κ is defined as the ratio of the energy in the specular line-of-sight path to the energy in the diffused components. Because of the line-of-sight component, the Ricean fading distribution is less “random” and has a lighter tail than the Rayleigh distribution with the same average channel gain. As a consequence, it can be seen from Figure 1.1 that the multiuser diversity gain is significantly smaller in the Ricean case compared to the Rayleigh case.

1.2 Multiuser versus Classical Diversity

We have called the above explained phenomenon multiuser diversity. Like the classical diversity techniques such as temporal, spatial and frequency

Fig. 1.1. Multiuser diversity gain for Rayleigh and Ricean fading channels ($\kappa = 5$); $\text{SNR} = P/N_0 = 0$ dB

diversity, multiuser diversity also arises from the existence of independently faded signal paths, in this case from the multiple users in the network. However, there are several important differences. First, the main objective of the classical diversity techniques is to improve the *reliability* of communication in slow fading channels; in contrast, the role of multiuser diversity is to increase the total throughput over fast fading channels. Under the sum capacity-achieving strategy, a user has no guarantee of a high rate in any particular slow fading state; only by averaging over the variations of the channel a high long-term average throughput is attained. Second, while the diversity techniques are designed to *counteract* the adverse effect of fading, multiuser diversity improves system performance by *exploiting* channel fading: channel fluctuations due to fading ensure that with high probability there is a user with a channel strength much larger than the mean level; by

allocating all the system resources to that user, the benefit of this strong channel is fully capitalized. Third, while the classical diversity techniques pertain to a point-to-point link, the benefit of multiuser diversity is *system-wide*, across the users in the network. This aspect of multiuser diversity has ramifications on the implementation of multiuser diversity in a cellular system. We discuss this next.

1.3 Multiuser Diversity: System Aspects

The cellular system requirements to extract the multiuser diversity benefits are:

- the base station has access to channel quality measurements: in the downlink, we need each receiver to track its own channel SNR, through say a common downlink pilot, and feed back the instantaneous channel quality to the base station (assuming a frequency division duplex (FDD) system).
- the ability of the base station to schedule transmissions among the users as well as to adapt the data rate as a function of the instantaneous channel quality.

These features are already present in the designs of many third-generation systems. Nevertheless, in practice there are several considerations to take into account before realizing such gains. In particular, there are three main hurdles towards a system implementation of the multiuser diversity idea.

- (i) **Fairness and Delay:** To implement the idea of multiuser diversity in a real system, one is immediately confronted with two issues: fairness and delay. In the ideal situation when users' fading *statistics* are the same, the strategy of communicating with the user having the best channel, maximizes not only the total capacity of the system but also the throughput of individual users. In reality, the statistics are not symmetric; there are users who are closer to the base station with a better average SNR; there are users who are stationary and some that are moving; there are users who are in a rich scattering environment and some with no scatterers around them. Moreover, the strategy is only concerned with maximizing long-term average throughputs; in practice there are latency requirements, in which case the average throughputs over the delay time-scale is the performance metric of interest. The challenge is to address these issues while at the same time exploiting the multiuser diversity gain inherent in a system with users having independent, fluctuating channel conditions. There are

many schedulers that harness multiuser diversity while addressing the real-world fairness issues; Chapter 6 of [Tse and Viswanath, 2005] has a detailed study of one of them: the *proportional fair* scheduler.

- (ii) **Channel Measurement and Feedback:** One of the key system requirements to harness multiuser diversity is to have scheduling decisions by the base station be made as a function of the channel states of the users. In the downlink, the users have access to their channel states but need to feedback these values to the base station. Both, the error in channel state measurement and the delay in feeding it back, constitute a significant bottleneck in extracting the multiuser diversity gains.

The prediction error is due to two effects: the error in measuring the channel from the pilot and the delay in feeding back the information to the base station. In the downlink, the pilot is shared between many users and is strong; so, the measurement error is quite small and the prediction error is mainly due to the feedback delay.

One remedy to reduce the feedback delay is to shrink the size of the scheduling time slot. However, this increases the feedback frequency in the uplink and thus, increases the system overhead. There are ways to reduce this feedback and Chapter 6 of [Tse and Viswanath, 2005] studies some of these techniques.

- (iii) **Slow and Limited Fluctuations:** We have observed that the multiuser diversity gains depend on the distribution of channel fluctuations. In particular, larger and faster variations in a channel are preferred over slow ones. However, there may be a line-of-sight path and little scattering in the environment, and hence the dynamic range of channel fluctuations may be small. Further, the channel may fade very slowly compared to the delay constraints of the application so that transmissions cannot wait until the channel reaches its peak. Effectively, the dynamic range of channel fluctuations is small within the time scale of interest. Both are important sources of hindrance to implementing multiuser diversity in a real system. We will next see a simple and practical scheme using an antenna array at the base station that creates fast and large channel fluctuations even when the channel is originally slow fading with a small range of fluctuation.

1.4 Opportunistic Beamforming using Dumb Antennas

The amount of multiuser diversity depends on the rate and dynamic range of channel fluctuations. In environments where the channel fluctuations are

Fig. 1.2. Same signal is transmitted over the two antennas with time varying phase and powers.

small, a natural idea comes to mind: why not amplify the multiuser diversity gain by *inducing* faster and larger fluctuations? Focusing on the downlink, we describe a technique that does this using multiple transmit antennas at the base station as illustrated in Figure 1.2.

Consider a system with n_t transmit antennas at the base station. Let $h_{lk}[m]$ be the complex channel gain from antenna l to the k^{th} user in time m . In time m , the same symbol $x[m]$ is transmitted from all of the antennas except that it is multiplied by a complex number $\sqrt{\alpha_l[m]}e^{j\theta_l[m]}$ at antenna l , for $l = 1 \dots n_t$, such that $\sum_{l=1}^{n_t} \alpha_l[m] = 1$, preserving the total transmit power. The received signal at user k (see the basic downlink fading channel model in (1.1) for comparison) is given by:

$$y_k[m] = \left(\sum_{l=1}^{n_t} \sqrt{\alpha_l[m]} e^{j\theta_l[m]} h_{lk}[m] \right) x[m] + w_k[m]. \quad (1.2)$$

In vector form, the scheme transmits $\mathbf{q}[m]x[m]$ at time m , where

$$\mathbf{q}[m] := \begin{bmatrix} \sqrt{\alpha_1[m]}e^{j\theta_1[m]} \\ \vdots \\ \sqrt{\alpha_{n_t}[m]}e^{j\theta_{n_t}[m]} \end{bmatrix} \quad (1.3)$$

is a unit vector and

$$y_k[m] = (\mathbf{h}_k[m]^* \mathbf{q}[m]) x[m] + w_k[m] \quad (1.4)$$

where $\mathbf{h}_k[m]^* := (h_{1k}[m], \dots, h_{n_t,k}[m])$ is the channel vector from the transmit antenna array to user k .

The overall channel gain seen by user k is now

$$\mathbf{h}_k[m]^* \mathbf{q}[m] = \sum_{l=1}^{n_t} \sqrt{\alpha_l[m]} e^{j\theta_l[m]} h_{lk}[m]. \quad (1.5)$$

The $\alpha_l[m]$'s denote the fractions of power allocated to each of the transmit antennas, and the $\theta_l[m]$'s denote the phase shifts applied at each antenna to the signal. By varying these quantities over time ($\alpha_l[m]$'s from 0 to 1 and $\theta_l[m]$'s from 0 to 2π), the antennas transmit signals in a time-varying direction, and fluctuations in the overall channel can be induced even if the physical channel gains $\{h_{lk}[m]\}$ have very little fluctuations (Figure 1.3).

As in the single transmit antenna system, each user k feeds back the overall received SNR of its own channel, $|\mathbf{h}_k[m]^* \mathbf{q}[m]|^2 / N_0$, to the base station (or equivalently the data rate that the channel can currently support) and the base station schedules transmissions to users accordingly. There is no need to measure the individual channel gains $h_{lk}[m]$ (phase or magnitude); in fact, the existence of multiple transmit antennas is completely transparent to the users. Thus only a single pilot signal is needed for channel measurement (as opposed to a pilot to measure each antenna gain). The pilot symbols are repeated at each transmit antenna, exactly like the data symbols.

The rate of variation of $\{\alpha_l[m]\}$ and $\{\theta_l[m]\}$ in time (or, equivalently, of the transmit direction $\mathbf{q}[m]$) is a design parameter of the system. We would like it to be as fast as possible to provide full channel fluctuations within the latency time scale of interest. On the other hand, there is a practical limitation to how fast this can be. The variation should be slow enough and should happen at a time scale that allows the channel to be reliably estimated by the users and the SNR fed back. Further, the variation should be slow enough to ensure that the channel seen by a user does not change abruptly and thus maintains stability of the channel tracking loop.

Fig. 1.3. Pictorial representation of the slow fading channels of two users before (above) and after (below) applying opportunistic beamforming.

1.4.1 Slow Fading: Opportunistic Beamforming

To get some insight into the performance of this scheme, consider the case of slow fading where the channel gain vector of each user k remains constant, i.e., $\mathbf{h}_k[m] = \mathbf{h}_k$, for all m . (In practice, this means: for all m over the latency time-scale of interest.) The received SNR for this user would have remained constant if only one antenna were used. If all users in the system experience such slow fading, no multiuser diversity gain can be exploited. Under the proposed scheme, on the other hand, the overall channel gain

$\mathbf{h}_k[m]^* \mathbf{q}[m]$ for each user k varies in time and provides an opportunity to exploit multiuser diversity.

Let us focus on a particular user k . Now if $\mathbf{q}[m]$ varies across all directions, the amplitude squared of the channel $|\mathbf{h}_k[m]^* \mathbf{q}[m]|^2$ seen by user k varies from 0 to $\sum_{l=1}^{n_t} |h_{lk}|^2$. The peak value occurs when the transmission is aligned along the direction of the channel of user k , i.e., $\mathbf{q}[m] = \mathbf{h}_k / \|\mathbf{h}_k\|$. The power and phase values are then in the *beamforming configuration* :

$$\begin{aligned} \alpha_l &= \frac{|h_{lk}|^2}{\sum_{j=1}^{n_t} |h_{jk}|^2}, & l = 1, \dots, n_t, \\ \theta_l &= -\arg(h_{lk}), & l = 1, \dots, n_t. \end{aligned}$$

To be able to beamform to a particular user, the base station needs to know individual channel amplitude and phase responses from all the antennas, which requires much more information to feedback than just the overall SNR. However, if there are many users in the system, a good multiuser diversity harnessing scheduler (the proportional fair algorithm is one such) will schedule transmission to a user only when its overall channel SNR is near its peak. Thus, it is plausible that in a slow fading environment, the technique can approach the performance of coherent beamforming but with only overall SNR feedback. In this context, the technique can be interpreted as *opportunistic beamforming*: by varying the phases and powers allocated to the transmit antennas, a beam is randomly swept and at any time transmission is scheduled to the user which is currently closest to the beam.

1.4.2 Fast Fading: Increasing Channel Fluctuations

We see that opportunistic beamforming can significantly improve performance in slow fading environments by adding fast time-scale fluctuations on the overall channel quality. The rate of channel fluctuation is artificially sped up. Can opportunistic beamforming help if the underlying channel variations are already fast (fast compared to the latency time-scale)?

The long term throughput under fast fading depends only on the stationary distribution of the channel gains. The impact of opportunistic beamforming in the fast fading scenario then depends on how the stationary distributions of the overall channel gains can be modified by power and phase randomization. Intuitively, better multiuser diversity gain can be exploited if the dynamic range of the distribution of h_k can be increased, so that the maximum SNRs can be larger. We consider two examples of common fading models.

- *Independent Rayleigh fading*: In this model, appropriate for an environment where there is full scattering and the transmit antennas are spaced sufficiently apart, the channel gains $h_{1k}[m], \dots, h_{n_t k}[m]$ are i.i.d. $\mathcal{CN}$ random variables. In this case, the channel vector $\mathbf{h}_k[m]$ is isotropically distributed, and $\mathbf{h}_k[m]^* \mathbf{q}[m]$ is circularly symmetric Gaussian for any choice of $\mathbf{q}[m]$; moreover the overall gains are independent across the users. Hence, the stationary statistics of the channel are *identical* to the original situation with one transmit antenna. Thus, in an independent fast Rayleigh fading environment, the opportunistic beamforming technique does not provide any performance gain.
- *Independent Ricean fading*: In contrast to the Rayleigh fading case, opportunistic beamforming has a significant impact in a Ricean environment, particularly when the κ -factor is large. In this case, the scheme can significantly increase the dynamic range of the fluctuations. This is because the fluctuations in the underlying Ricean fading process come from the diffused component, while with randomization of phase and powers, the fluctuations are from the coherent addition and cancellation of the direct path components in the signals from the different transmit antennas, in addition to the fluctuation of the diffused components. If the direct path is much stronger than the diffused part (large κ values), then much larger fluctuations can be created with this technique.

This intuition is substantiated in Figure 1.4, which plots the total throughput with the proportional fair algorithm (with large latency time scale) for Ricean fading with $\kappa = 10$. We see that there is a considerable improvement in performance going from the single transmit antenna case to dual transmit antennas with opportunistic beamforming. For comparison, we also plot the analogous curves for pure Rayleigh fading; as expected, there is no improvement in performance in this case. Figure 1.5 compares the stationary distributions of the overall channel gain $|\mathbf{h}_k[m]^* \mathbf{q}[m]|$ in the single-antenna and dual-antenna cases; one can see the increase in dynamic range due to opportunistic beamforming.

1.5 Antennas: Dumb, Smart and Smarter

It is insightful to compare the opportunistic beamforming technique with the two other important point-to-point transmit antenna techniques:

- *space-time codes* like the Alamouti scheme. They are primarily used to increase the diversity in slow fading point-to-point links.

Fig. 1.4. Total throughput as a function of the number of users under Ricean fast fading, with and without opportunistic beamforming. The power allocation $\alpha_l[m]$'s are uniformly distributed in $[0, 1]$ and the phases $\theta_l[m]$'s uniform in $[0, 2\pi]$.

- *transmit beamforming*. In addition to providing diversity, a power gain is also obtained through the coherent addition of signals at the users.

The three techniques have different system requirements. Coherent space-time codes like the Alamouti scheme require the users to track all the *individual* channel gains (amplitude and phase) from the transmit antennas. This requires separate pilot symbols on each of the transmit antennas. Transmit beamforming has an even stronger requirement that the channel should be known at the transmitter. In an FDD system, this means feedback of the individual channel gains (amplitude and phase). In contrast to these two techniques, the opportunistic beamforming scheme requires no knowledge of the individual channel gains, neither at the users nor at the transmitter. In fact, the users are *completely ignorant* of the fact that there are multiple

Fig. 1.5. Comparison of the distribution of the overall channel gain with and without opportunistic beamforming using two transmit antennas, Ricean fading.

transmit antennas and the receiver is identical to that in the single transmit antenna case. Thus, they can be termed *dumb antennas*. Opportunistic beamforming does rely on multiuser diversity scheduling, which requires the feedback of the overall SNR of each user. However, this only needs a *single* pilot to measure the overall channel.

What is the performance of these techniques when used in the downlink? In a slow fading environment, we have already remarked that opportunistic beamforming approaches the performance of transmit beamforming when there are many users in the system. On the other hand, space-time codes do not perform as well as transmit beamforming since they do not capture the array power gain. This means, for example, using the Alamouti scheme on dual transmit antennas in the downlink is 3 dB worse than using opportunistic beamforming combined with multiuser diversity scheduling when

there are many users in the system. Thus, dumb antennas together with smart scheduling can surpass the performance of smart space-time codes and approach that of the even smarter transmit beamforming.

How about in a fast Rayleigh fading environment? In this case, we have observed that dumb antennas have no effect on the overall channel as the full multiuser diversity gain has already been realized. Space-time codes, on the other hand, *increase* the diversity of the point-to-point links and consequently *decrease* the channel fluctuations and hence the multiuser diversity gain. Thus, the use of space-time codes as a point-to-point technology in a multiuser downlink with rate control and scheduling can actually be *harmful*, in the sense that even the naturally present multiuser diversity is removed. The performance impact of using transmit beamforming is not so clear: on the one hand it reduces the channel fluctuation and hence the multiuser diversity gain, but on the other hand it provides an array power gain. However, in an FDD system the fast fading channel may make it very difficult to feed back so much information to enable coherent beamforming.

1.6 Multiuser Diversity in Multi-cell Systems

So far we have considered a single-cell scenario, where the noise is assumed to be white Gaussian. For wide band cellular systems with full frequency reuse, it is important to consider the effect of inter-cell interference on the performance of the system, particularly in interference-limited scenarios. In a cellular system, this effect is captured by measuring the channel quality of a user by the SINR, signal-to-interference-plus-noise ratio. In a fading environment, the energies in both the received signal and the received interference fluctuate over time. Since the multiuser diversity scheduling algorithm allocates resources based on the channel SINR (which depends on both the channel amplitude and the amplitude of the interference), it automatically exploits both the fluctuations in the energy of the received signal as well as that of the interference: the algorithm tries to schedule resource to a user whose instantaneous channel is good *and* the interference is weak. Thus, multiuser diversity naturally takes advantage of the time-varying interference to increase the spatial reuse of the network.

From this point of view, amplitude and phase randomization at the base station transmit antennas plays an additional role: it increases not only the amount of fluctuations of the received signal to the intended users *within* the cells, it also increases the fluctuations of the interference the base station causes in *adjacent* cells. Hence, opportunistic beamforming has a dual benefit in an interference-limited cellular system. In fact, opportunistic beam-

forming performs *opportunistic nulling* simultaneously: while randomization of amplitude and phase in the transmitted signals from the antennas allows near coherent beamforming to some user within the cell, it will create near nulls at some other user in adjacent cells. This in effect allows *interference avoidance* for that user if it is currently being scheduled.

Let us focus on the downlink and slow flat fading scenario to get some insight into the performance gain from opportunistic beamforming and nulling. Under amplitude and phase randomization at all base stations, the received signal of a typical user that is interfered by J adjacent base stations is given by

$$y[m] = (\mathbf{h}^* \mathbf{q}[m]) x[m] + \sum_{j=1}^J (\mathbf{g}_j^* \mathbf{q}_j[m]) u_j[m] + z[m]. \quad (1.6)$$

Here, $x[m]$, $\mathbf{h}$, $\mathbf{q}[m]$ are respectively the signal, channel vector and random transmit direction from the base station of interest; $u_j[m]$, $\mathbf{g}_j$, $\mathbf{q}_j[m]$ are respectively the interfering signal, channel vector and random transmit direction from the j^{th} base station. All base stations have the same transmit power, P , and n_t transmit antennas and are performing amplitude and phase randomization independently.

By averaging over the signal $x[m]$ and the interference $u_j[m]$'s, the (time-varying) SINR of the user k can be computed to be

$$\text{SINR}_k[m] = \frac{P |\mathbf{h}^* \mathbf{q}[m]|^2}{P \sum_{j=1}^J |\mathbf{g}_j^* \mathbf{q}_j[m]|^2 + N_0}. \quad (1.7)$$

As the random transmit directions $\mathbf{q}[m]$, $\mathbf{q}_j[m]$ vary, the overall SINR changes over time. This is due to the variations of the overall gain from the base station of interest as well as those from the interfering base stations. The SINR is high when $\mathbf{q}[m]$ is closely aligned to the channel vector $\mathbf{h}$, and/or for many j 's, $\mathbf{q}_j[m]$ is nearly orthogonal to $\mathbf{g}_j$, i.e., the user is near a null of the interference pattern from the j^{th} base station. In a system with many other users, the proportional fair scheduler will serve this user while its SINR is at its peak $P \|\mathbf{h}\|^2 / N_0$, i.e., when the received signal is the strongest and the interference is completely nulled out. Thus, the opportunistic nulling and beamforming technique has the potential of shifting a user from a low SINR, interference-limited regime to a high SINR, noise-limited regime.

1.7 A Concluding System View

A new design principle for wireless systems can now be seen through the lens of multiuser diversity. In the classical view point, the design techniques

center on making the *individual* point-to-point links as close to AWGN channels as possible, with a reliable channel quality that is constant over time. This is accomplished by *channel averaging*, and includes the use of diversity techniques such as multipath combining, time interleaving and antenna diversity that attempt to keep the channel fading constant in time, as well as interference management techniques such as interference averaging by means of spreading.

However, if one shifts from the view of the wireless system as a set of point-to-point links to the view of a system with multiple users sharing the same resources (spectrum and time), then quite a different design objective suggests itself. Indeed, the results in this chapter suggest that one should instead try to *exploit* the channel fluctuations. This is done through an appropriate scheduling algorithm that “rides the peaks”, i.e., each user is scheduled when it has a very strong channel, while taking into account real world traffic constraints such as delay and fairness. The technique of dumb antennas goes one step further by *creating* variations when there are none. This is accomplished by varying the strengths of *both* the signal and the interference that a user receives through opportunistic beamforming and nulling.

The viability of the opportunistic communication scheme depends on traffic that has some tolerance to scheduling delays. On the other hand, there are some forms of traffic that are not so flexible. The functioning of the wireless systems is supported by the overhead control channels which are “circuit-switched” and hence have very tight latency requirements, unlike data which have the flexibility to allow dynamic scheduling. From the perspective of these signals, it is preferable that the channel remained unfaded; a requirement that is contradictory to our scheduler-oriented observation that we would prefer the channel to have fast and large variations.

This issue suggests the following design perspective: separate very-low latency signals (such as control signals) from flexible latency data. One way to achieve this separation is to split the bandwidth into two parts. One part is made as flat as possible (by spreading over this part of the bandwidth, say) and is used to transmit flows with very low latency requirements. The performance metric here is to make the channel as reliable as possible (equivalently keeping the probability of outage low) for some fixed data rate. The second part uses opportunistic beamforming to induce large and fast channel fluctuations and a scheduler to harness the multiuser diversity gains. The performance metric on this part is to maximize the multiuser diversity gain.

The gains of the opportunistic beamforming and nulling depend on the probability that the received signal is near beamformed *and* all the interfer-

ence is near null. In the interference-limited regime and when $P/N_0 \gg 1$, the performance depends mainly on the probability of the latter event. In the downlink, this probability is large since there are only one or two base stations contributing most of the interference. The uplink poses a contrasting picture: there is interference from many mobiles allowing interference averaging. Now the probability that the *total* interference is near null is much smaller. Interference averaging, which is one of the principle design features of wideband full reuse systems is actually unfavorable for the opportunistic scheme described here, since it reduces the likelihood of the nulling of the interference and hence the likelihood of the peaks of the SINR.

In a typical cell, there will be a distribution of users, some closer to the base station and some closer to the cell boundaries. Users close to the base station are at high SINR and are noise-limited; the contribution of the inter-cell interference is relatively small. These users benefit mainly from opportunistic beamforming. Users close to the cell boundaries, on the other hand, are at low SINR and are interference-limited; the average interference power can be much larger than the background noise. These users benefit both from opportunistic beamforming and from opportunistic nulling of inter-cell interference. Thus, the cell-edge users benefit more in this system than users in the interior. This is rather desirable from a system fairness point-of-view, as the cell-edge users tend to have poorer service. This feature is particularly important for a system without soft handoff (which is difficult to implement in a packet data scheduling system). To maximize the opportunistic nulling benefits, the transmit power at the base station should be set as large as possible, subject to regulatory and hardware constraints.

We have seen the multiuser diversity as primarily a form of power gain. The opportunistic beamforming technique of using an array of multiple transmit antennas has approximately an n_t fold improvement in received SNR to a user in a slow fading environment, as compared to the single-antenna case. With an array of n_r *receive* antennas at each mobile (and say a single transmit antenna at the base station), the received SNR of any user gets an n_r fold improvement as compared to a single receive antenna; this gain is realized by *receiver beamforming*. This operation is easy to accomplish since the mobile has full channel information at each of the antenna elements. Hence the gains of opportunistic beamforming are about the same order as that of installing sets of receive antenna arrays at *each* of the mobiles.

Thus, for a system designer, the opportunistic beamforming technique provides a compelling case for implementation, particularly in view of the

constraints of space and cost of installing multiple antennas on *each* mobile device. Further, this technique neither needs any extra processing on part of any user, nor any updates to an existing air-link interface standard. In other words, the mobile receiver can be completely ignorant to the use or non-use of this technique. This means that it does not have to be “designed in” (by appropriate inclusions in the air interface standard and the receiver design) and can be added/removed at any time. This is one of the important benefits of this technique from an overall system design point of view.

In traditional cellular wireless systems, the cell is sectorized to allow better focusing of the power transmitted from the antennas and also to reduce the interference seen by mobile users from transmissions of the same base station but intended for users in different sectors. This technique is particularly gainful in scenarios when the base station is located at a fairly large height and thus there is limited scattering around the base station. In contrast, in systems with far denser deployment of base stations (a strategy that can be expected to be a good one for wireless systems aiming to provide mobile, broadband data services), it is unreasonable to stipulate that the base stations be located high above the ground so that the local scattering (around the base station) is minimal. In an urban environment, there is substantial local scattering around a base station and the gains of sectorization are minimal; users in a sector also see interference from the same base station (due to the local scattering) intended for another sector. The opportunistic beamforming scheme can be thought of as sweeping a random beam and scheduling transmissions to users when they are beamformed. Thus, the gains of sectorization are automatically realized. We conclude that the opportunistic beamforming technique is particularly suited to harness sectorization gains even in low height base stations with plenty of local scattering. In a cellular system, the opportunistic beamforming scheme also obtains the gains of nulling, a gain traditionally obtained by coordinated transmissions from neighboring base stations in a full frequency reuse system or by appropriately designing the frequency reuse pattern.

1.8 Acknowledgements

The research that led to the material in this chapter was supported in part by the National Science Foundation under grant CCR 0312413 and by the Motorola Center for Communication. The material in this chapter is adapted from [Viswanath et al., 2002] and Chapter 6 of [Tse and Viswanath, 2005].

References

- D. Tse and P. Viswanath. *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
- P. Viswanath, D. Tse, and R. Laroia. Opportunistic beamforming using dumb antennas. *IEEE Transactions on Information Theory*, 48(6):1277–1294, June 2002.