Overcoming Interference in Spatial Multiplexing MIMO Cellular Networks

Jeffrey G. Andrews, Wan Choi, and Robert W. Heath Jr.

Submitted to *IEEE Wireless Communications Magazine*

Last Modified: October 20, 2006

Abstract

Multi-antenna transmission and reception (known as MIMO) is widely touted as the key technology for enabling wireless broadband services, whose widespread success will require ten times higher spectral efficiency than current cellular systems, at ten times lower cost per bit. Spectrally efficient, inexpensive cellular systems are by definition densely populated and interference-limited. But spatial multiplexing MIMO systems – whose principal merit is a supposed dramatic increase in spectral efficiency – lose much of their effectiveness in high levels of interference. This paper overviews several approaches for handling interference in multicell MIMO systems. The discussion is applicable to any multi-antenna cellular network including 802.16e/WiMAX, 3GPP (HSDPA and 3GPP LTE) and 3GPP2 (1xEVDO). We argue that many of the traditional interference management techniques have limited usefulness (or are even counterproductive) when viewed in concert with MIMO. The problem of interference in MIMO systems is too large in scope to be handled with a single technique: in practice a combination of complementary countermeasures will be needed. We overview emerging system-level interference-reducing strategies based on cooperation, which will be important for overcoming interference in future spatial multiplexing cellular systems.

The authors are with the Wireless Networking and Communications Group, Department of Electrical and Computer Engineering, The University of Texas at Austin, 1 University Station C0803, Austin, TX 78712, USA. Email:{jandrews, rheath}@ece.utexas.edu, wchoi@mail.utexas.edu, Phone:1-512-471-0536.

I. The Promise of MIMO

The idea of using multiple receive and multiple transmit antennas has emerged as one of the most significant technical breakthroughs in modern wireless communications. Theoretical studies and initial prototyping of these multiple-input multiple-output (MIMO) systems have shown order of magnitude spectral efficiency improvements in point-to-point communication. As a result, MIMO is considered a key technology for improving the throughput of future wireless broadband data systems, which presently are mired at data rates far below their wired and wireless LAN counterparts. For example, contemporary[1] "broadband" cellular systems like HSDPA and 1xEV-DO achieve physical layer data rates around 0.3 (uplink) – 0.8 (downlink) bps/Hz/sector, whereas modern wireless LANs achieve around 4-5 bps/Hz in both directions.

The multidimensional MIMO channel can be exploited to increase diversity or to provide parallel spatial channels, which is known as spatial multiplexing. Transmit and receive diversity is generally considered a more conservative approach; a well-known example is space-time codes [1], [2], which have found wide adoption in 3G CDMA cellular systems [3] as well as in 802.16/WiMAX. Assume the transmitter has M_t antennas and the receiver M_r. Diversity increases the robustness of the system by eliminating fades; channel-aware diversity[2] also raises the average received signal to noise ratio (SNR) in proportion to M_t and/or M_r. If the channel is not known, then the diversity hardens the SNR to the mean SNR rather than increasing it – that is, it converts a fading channel to a non-fading channel – as the number of antennas grows large. Even if the average SNR increases linearly with the number of antennas, the capacity growth is *logarithmic*, easily verified with Shannon's formula $C = B\log_2(1 + SNR)$. Spatial multiplexing, on the other hand, divides the incoming data into multiple parallel substreams and transmits each on a different spatial dimension (e.g. a different antenna). As long as there at

[1] As of late 2006.

[2] For example, coherent combining at the receiver, or channel-aware precoding at the transmitter.

least as many (sufficiently spaced) receive antennas as transmitted streams, spatial multiplexing increases the capacity *linearly* with the number of streams [4].

In this paper, we focus primarily on the spatial multiplexing aspect of MIMO since it is the most aggressive approach for increasing the link capacity, and since transmit diversity schemes are already widely implemented[3]. Because high data rates are particularly interesting for the downlink, it can be assumed that the number of transmit antennas M_t will be larger than the number of receive antennas M_r, due to space and cost restrictions on the mobile unit. It is reasonable to expect that M_r data streams will be transmitted leaving $M_t - M_r$ degrees of freedom for achieving transmit diversity using antenna subset selection [5], [6] or transmit precoding with limited feedback [7] for example.

The main goals of the paper are as follows. *First*, we convey the severity of the interference problem in a cellular scenario, and explain why a decade after its invention, spatial multiplexing has still not been used in cellular networks. *Second*, we describe recently developed techniques for combating interference in MIMO systems. While some of those techniques are useful, most of them directly compete with the main goal of spatial multiplexing, which is high capacity. Some of the techniques are in fact counterproductive. *Third*, we describe the recent trend towards strategic interference-reduction based on base station cooperation and distributed antennas (which is also a form of cooperation), and argue that such techniques hold the key to making MIMO realizable in future cellular-based communication systems.

II. Spatial Multiplexing in Cellular Systems

The vast majority of academic and even industrial research has focused on the point-to-point MIMO model (well summarized in [8]), which ignores nearby competing interference sources. This is reasonable in certain applications. The IEEE 802.11n standard, which is a high data rate MIMO extension of IEEE 802.11g, uses the medium access control protocol to ensure that

[3]In this paper the term "MIMO", without qualification, should be interpreted to mean spatial multiplexing.

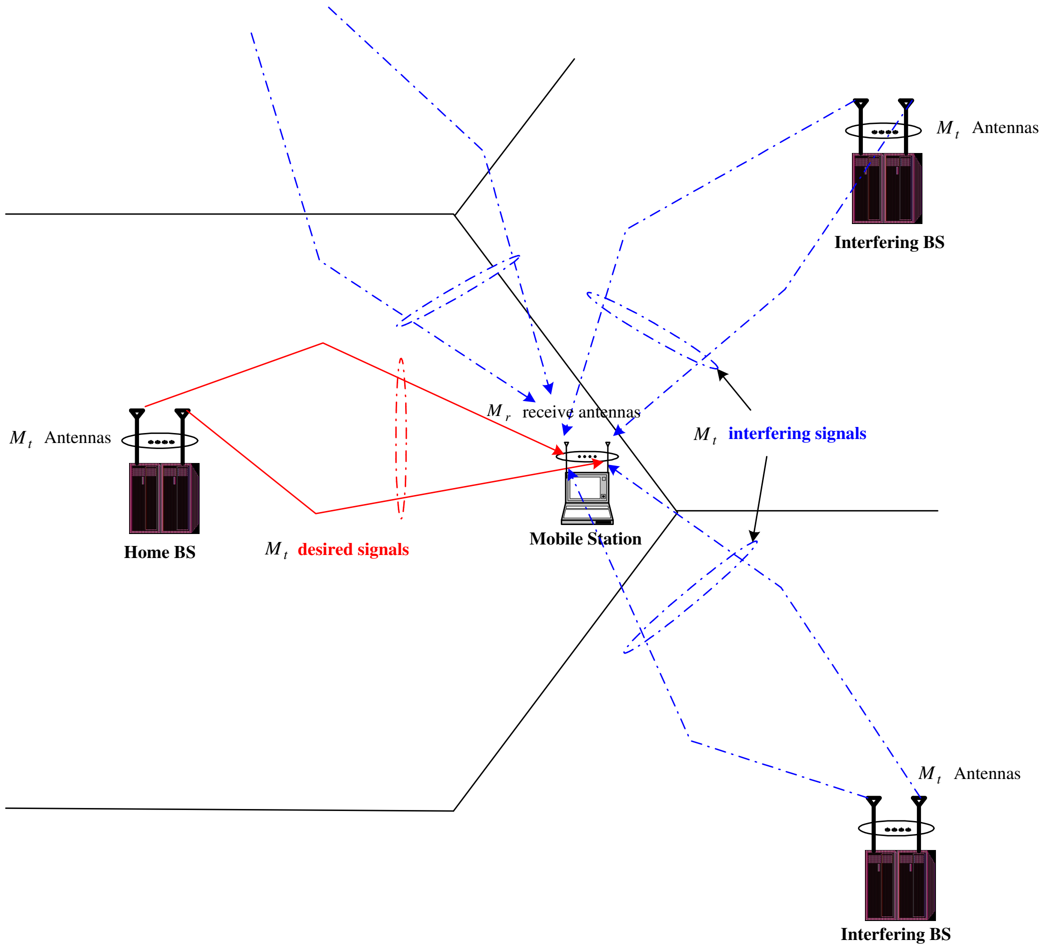

Fig. 1. Other-cell interference in MIMO cellular systems

its short range links do not suffer from interference, which is reasonable in all but the densest deployments. It is quite a different matter though to apply MIMO successfully in cellular systems, since due to cell planning and coverage considerations, cellular systems are interference limited. Despite these challenges, MIMO is being widely considered for next generation cellular systems such as IEEE802.16/WiMAX and the 3GPP Long Term Evolution (LTE).

A. The Other Cell Interference (OCI) Problem in Cellular MIMO Systems

All well-designed cellular systems are by nature interference-limited: if they were not, it would be possible to increase the spectral efficiency by lowering the frequency reuse or increasing the average loading per cell. Note that this paper is focused on *other-cell interference* (OCI), as opposed to the *self-cell interference* from other users in the cell (as is typical in an uplink CDMA system) or the *co-antenna interference* resulting between the spatial multiplexing data streams of a single user. Methods for the latter two types of interference are well-understood, and users in each cell of wireless broadband networks are generally orthogonal: through TDMA and OFDMA in 802.16 systems, or through orthogonal CDMA (Walsh codes) coupled with a chip-level equalizer and TDMA in the 3GPP systems.

The downlink of a cellular system is expected to be the most profitable and viable for MIMO communication. Unfortunately, as illustrated in Figure 1, there will be NM_t interfering signals if there are N non-negligible neighboring base stations. The total interference power does not necessarily increase since the typical assumption in MIMO systems is that the transmit power per antenna is reduced by M_t, so that the total transmitted power is the same as in the non-MIMO setting. The number of interfering signals, though, does increase. From straightforward linear algebra arguments, MIMO receivers are able to decode parallel data streams by suppressing the spatial interference between the signals sent from the M_t transmit antennas (using linear signal processing techniques), as long as the number of receive antennas $M_r \geq M_t$. Therefore, an interference-dominated MIMO system requires $M_r \geq (N+1)M_t$ receive antennas in order to fully suppress OCI (with a linear receiver), which is out of the question on small handsets. In short, the number of interfering streams does not affect the power of the interference, but it has a substantial affect on its statistical distribution.

As the number of interfering streams increases, it is not typically possible to suppress all the OCI with spatial signal processing, and instead the interference is generally treated as noise. As

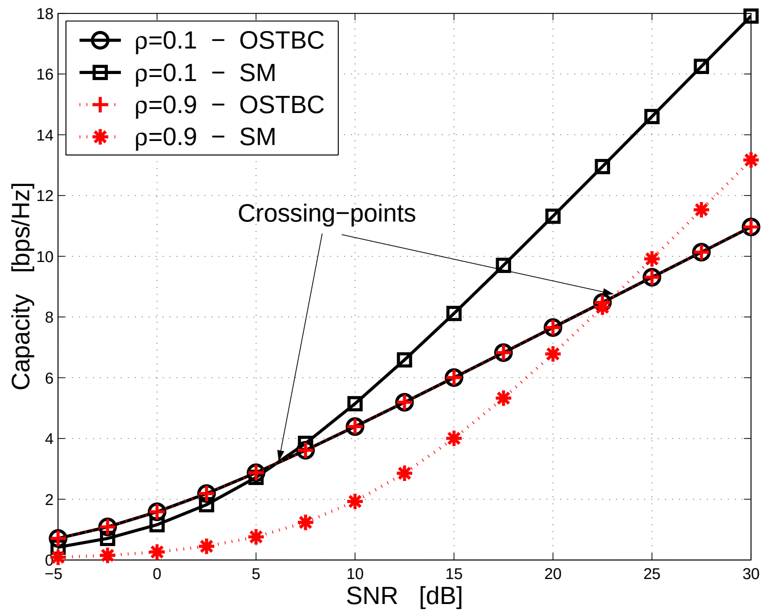

Fig. 2. Capacity of spatial multiplexing and OSTBCs vs. SINR for 2×2 MIMO systems with spatial correlation of $\rho = \{0.1, 0.9\}$ [11].

the number of interfering sources becomes large, the interference becomes increasingly Gaussian due to the Central Limit Theorem, which is worst case according to information theory. In view of this, it is not surprising that recent research on spatial multiplexing in cellular systems has reached the common conclusion that *adding active transmit antennas or data streams at each base station can actually decrease the throughput at low SINR* due to the increased dimensionality of spatial interference [9], [10]. Instead, at low SINR, higher capacity is achieved by focusing power on the best eigenvalue(s) in the channel matrix.

B. The gulf between "theory" and theory

It is often joked that "in theory, theory and practice are the same". The large gap between theory and practice is often a result of compromising assumptions made in the development of a theory. For example, consider the conventional wisdom that spatial multiplexing has a fundamental capacity advantage relative to transmit diversity. In many practical situations, even

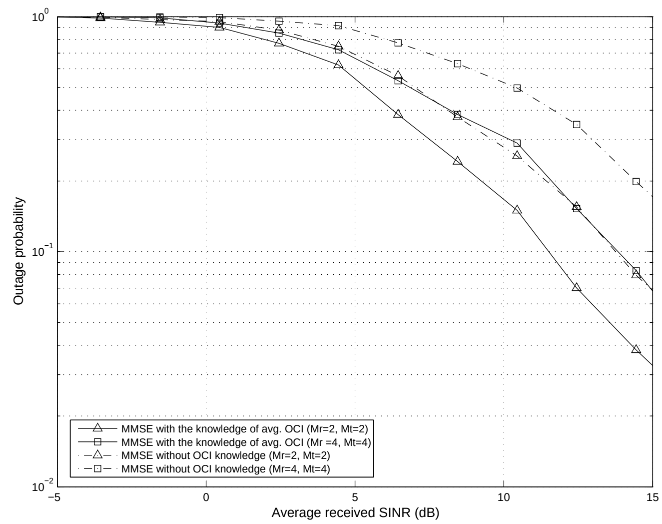

Fig. 3. Outage probability of a cellular MIMO system with target Eb/No = 1.5dB. The shadowing is assumed to have $\sigma = 8dB$ for interference and 2dB for the desired user. Note that even with this very low target Eb/No, the outage probability is very high.

without considering OCI, diversity outperforms spatial multiplexing in terms of both data rate and reliability. For example, both [8], [11] show that at low SINR, space-time block codes achieve higher capacity and throughout than spatial multiplexing, especially if there is non-trivial spatial correlation. For example, Fig. 2 shows that for SINRs below 5 dB in highly correlated channels, or even up to SINRs approaching 20 dB in channels with low correlation, simple orthogonal space time block codes (OSTBCs) achieve higher capacity than spatial multiplexing.

The lesson is not that theory is wrong, but rather, that the assumptions made are quite important. Inappropriate assumptions – e.g. high SINR, uncorrelated antenna arrays, neglect of user geometry in the cells – can lead to conclusions that are quite misleading for the majority of users.

Evidence against neglecting interference when considering MIMO techniques is growing. Recent Bell Labs research [12] has pointed out that the typical SINR operating point is about 2dB in modern cellular systems, at which Shannon capacity has been approached within reasonable limits. Making matters worse, at least 10% of the users in the cell have SINRs that are considerably lower than 0 dB in an aggressive time and frequency reuse pattern. Nevertheless, most academic research has continued to ignore other-cell interference because such interference is difficult to handle analytically, and the results for MIMO with OCI are often disappointing since the multiplicative gains in capacity predicted for spatial multiplexing only occur at high SINR. This fact, combined with the lack of additional degrees of freedom to deal with interference spells disaster for spatial multiplexing in cellular systems.

To make this argument more concretely, consider Fig. 3. Here, the outage probability for a conservative *target* Eb/No of 1.5 dB is plotted for 2×2 and 4×4 MIMO with a linear MMSE receiver that uses the average OCI level to balance OCI-noise enhancement[4] with spatial interference suppression. As can be seen, even with an optimistic *average* received SINR (say 10 dB), the outage performance is very poor for spatial multiplexing due to the difficulty in balancing the spatial interference and other-cell interference. Furthermore, the outage probability increases as the number of antennas increases, due to the increased spatial dimensionality. In summary, there is a serious conflict between spatial multiplexing and heavily loaded cellular systems. In the rest of the paper, we enumerate and discuss the prominent strategies for interference-limited multicell MIMO systems – some of which have been proposed very recently – and debate their relative merits.

III. Interference Mitigation Techniques for Wireless Broadband Networks

In this section, we outline possible OCI mitigation techniques in cellular MIMO systems and discuss their feasibility. The discussed techniques are summarized and compared in Table I.

[4] The MMSE performance using only thermal noise is shown for comparison - the associated loss is a few dB.

A. Conventional Approaches to Dealing with Interference in Cellular Systems

The problem of other-cell interference has existed in cellular systems for many years. Traditionally, there are several different approaches to interference management including frequency reuse, sectoring, and spread spectrum. These techniques can be applied to combat OCI in MIMO cellular systems but each has important drawbacks.

Frequency reuse has been adopted for OCI reduction in cellular systems owing to its simplicity and practicality. With frequency reuse, the available spectrum is divided among cells in a cluster and reused in subsequent clusters. Frequency reuse effectively reduces OCI by spacing the competing transmissions farther away, and particularly benefits users near the cell boundaries. This of course reduces the spectral efficiency since for a frequency reuse of f only $1/f$ of the available spectrum is used each cell. A typical value of f in systems without spread spectrum is around three. Universal frequency reuse ($f = 1$) is highly desirable since it simplifies the problem of cell planning and reduces the frequency reuse penalty, but this results in intolerably high interference for MIMO systems. Therefore, frequency reuse does not solve the MIMO OCI problem since it is unlikely that MIMO's gains of MIMO will outweigh the $1/f$ penalty without a very large number of antennas.

Cell sectoring is another technique used to reduce the average interference at the edge of a cell. The idea of sectoring is to employ directive antennas at the tower to restrict the radiation at the base station. Sectoring improves the average SINR and improves spectral efficiency since frequencies can be reused in each sector. Cell sectoring, which is highly effective in current cellular systems, has the potential drawback of reducing the multipath diversity in the channel, which in the absence of sufficient local scattering can reduce the benefits of MIMO. Small sectors also result in unbalanced numbers of users per sector, which limits the capacity per user in the highly populated sectors, while wasting capacity in the thinly populated sectors. Sixty degree sectors (6 per cell) are considered a practical maximum. Regardless, sectoring is more attractive than large frequency reuse since it does not entail as large of a reuse penalty.

Another effective approach in modern cellular systems is to attempt to average the interference levels using spread spectrum techniques: either by direct sequence CDMA, or by slow frequency hopping in GSM. In these systems OCI is controlled by adjusting the traffic load in each cell to be kept significantly below what would be acceptable in a single-cell system. This acceptable load drops further in a MIMO system due to its susceptibility to interference [13].

B. Advanced Receiver Techniques

MIMO receivers that are interference-aware can significantly attenuate the OCI and hence greatly improve system performance. Most of the multiuser receiver literature is concerned with *self-cell* interference, but as noted in Section II-A, this is not likely relevant to future wireless broadband systems. Receivers for dealing with co-antenna interference have been studied extensively in the MIMO literature including linear, near optimal, and optimal techniques. Incorporating OCI into the co-antenna interference receiver is one of the main challenges in MIMO receiver design. Therefore, we now overview advanced signal processing options available to the receiver to suppress *other-cell interference*, and explain the key challenges these techniques face if they are to contribute to the adoption of MIMO cellular systems.

1) (Near) Maximum Likelihood Multiuser Detection: If instantaneous information on the channels of interferers is available, maximal likelihood (ML) multiuser detection (MUD) is known to minimize the bit error probability in a multicellular MIMO system. However, not only is this instantaneous channel information difficult to attain for neighboring base stations, the complexity of such a receiver is prohibitive for a low-power mobile unit. For an $M_r \times M_r$ MIMO system using M-QAM, the complexity is on the order of M^{NM_r}, where N is the number of interferers [14]. Even for simple cases this is well beyond reasonable implementation complexity. An alternative is the sphere decoder, which searches for the optimal solution in a sphere around an estimate of the received codeword. Sphere decoders can deal with moderate NM_r but still fail for large numbers of antennas or interferers. Additionally, they have variable complexity making implementation difficult. VLSI implementations are under development [15] but their viability in power hungry mobiles is still open to debate.

2) The MMSE Receiver: Due to the complexity problems associated with the ML receiver, a natural approach is to consider a linear approximation to the ML receiver. The zero-forcing spatial receiver performs terribly due to its enhancement of other-cell interference, which suggests the need for an MMSE receiver that balances noise-enhancement with spatial interference suppression. Here, two classes of MMSE receivers can be considered: one that knows the interferer's channels (an MMSE multiuser detector) and one that only knows the average OCI (typical MMSE) [16]. As discussed previously with regards to Figure 3, the MMSE receiver with only average OCI knowledge, while significantly better than a zero-forcing receiver, does not perform particularly well due to the inevitable other-cell interference enhancement.

Naturally, the MMSE multiuser detector has superior performance since it is able to explicitly reject the other-cell interference. It suffers, however, from two important problems. First, as in the ML detector, a major difficulty arises in obtaining instantaneous channel knowledge for the interfering base stations. Second, because M_t antennas at each of N interfering base stations must be treated as M_t independent interferers (for a total of M_tN interferers), effectively suppressing all of this interference is very difficult with only $M_r \ll M_tN$ receive antennas. One simplification, which is being looked at by some service providers, is to cancel just the pilot and common channel signals of the strongest interfering base station, which also alleviates the channel estimation requirements. The network-wide capacity gains from doing this are predicted to be around 10-25% (the fraction of power that these overhead signals contribute), which is perhaps worthy of implementation but not a major leap forward.

3) Interference Cancelling Receivers: As observed in single antenna systems, nonlinear receivers often provide a desirable tradeoff between performance and complexity [17], relative to ML receivers at one extreme and linear receivers at the other. In multicell MIMO systems, group detection techniques have a natural appeal, in which information bits for each "group" (cell) are detected sequentially [14]. One of the most popular among various group detection techniques is *group decision feedback multiuser detector*, an extension from BLAST. This receiver detects one MIMO system at a time and then feeds the tentative decisions to other group detectors for interference cancellation. Although successive interference cancellation is asymptotically optimal under the assumption of perfect interference cancellation, it is very susceptible to inaccurate channel estimates. Further, interference canceling receivers are still likely too complex for low-power mobile units.

4) Comments on Channel Information for Interfering Base Stations: An important shortcoming of the above techniques is that most require (near-instantaneous) channel information for the neighboring base stations, which is traditionally not available. In cellular systems, mobile stations periodically monitor pilot channels of neighboring base stations to assist with handoff. These pilot signals could potentially be used to gather the required channel knowledge for the above multiuser receivers. Still, such schemes appear too complex for the downlink, where each mobile would have to frequently monitor neighboring base stations and process updates to a complicated multiuser detector every channel coherence time.

C. Advanced Transmitter Techniques

For many systems, in both theory and practice, transmitter adaptation has been shown to be a highly profitable means of exploiting channel conditions. A well-known example is adaptive modulation, which is adopted to achieve high data rates for users with good channels. Another method is the closed-loop diversity, which approximates the maximum ratio transmission beamforming vector. Transmitter-based techniques have the additional merit in a cellular downlink of transferring the complexity burden from the mobile unit to the base station, where higher complexity is more tolerable and additional antennas may be available. In this section, we consider methods of using only self-cell channel knowledge at the transmitter to help suppress other-cell interference using advanced signal processing techniques. Multicell transmitter cooperation is discussed in the next section.

1) Closed-Loop MIMO Diversity Schemes: Given $M_t > M_r$ antennas, it is possible to improve the diversity performance of spatial multiplexing by using channel state information at the transmitter. One example is antenna subset selection where the best M_r of M_t antennas are selected, but antenna selection is not particularly effective in wideband systems due to the plentiful frequency diversity [6]. Other closed-loop possibilities include eigenbeamforming and transmit precoding [18]. These schemes require knowledge of channel state information at the transmitter; efficient limited feedback strategies for this purpose have been developed recently [7]. Closed-loop diversity schemes provide a diversity improvement of $M_t - M_r + 1$ and additional array gain, which reduces the required transmit power for spatial multiplexing, and hence the interference caused to neighboring cells. On a system-wide level, these approaches have about the same impact on OCI as transmit diversity with power control.

2) Stream Control in MIMO Systems: One simple but effective means for reducing OCI is simply to spatially multiplex $M < \min(M_t, M_r)$ data streams. This concept is known as *stream control* when used to reduce interference [19] and as *multi-mode adaptation* when combined with transmit precoding as a means to achieve additional diversity with low complexity receivers (e.g. [20]). Although the nominal transmitted data rate of each user is reduced compared to sending the maximum possible number of streams, stream control can actually improve the overall system capacity in many cases. In fact, for low SINRs sending a single stream using transmit beamforming is usually optimal from a sum capacity perspective. Stream control does not confront the problem of OCI directly, rather it acquiesces by gracefully reducing the number of streams, thus achieving the double benefit of concentrating its own power on its best eigenvalue(s) while reducing the dimensionality of the interference it causes to others.

Stream control is typically implemented in conjunction with adaptive modulation on a per-user basis. Users that are close to the base station, for example, would naturally support more streams than users on the cell boundaries due to their higher SINR. Due to their proximity to the base station, the transmit power to those users can be reduced even when transmitting multiple streams, so that the interference generated to other cells is within limits. In summary, adaptive stream control is a possible method for managing OCI in a MIMO cellular system and will likely be a component of any holistic OCI management solution, but is not itself necessarily the OCI solution since it simply tries to avoids the problem.

3) MIMO Combined with Transmit Beamforming for Interference Reduction: Beamforming is a term applied to a large number of different techniques including eigenbeamforming and precoding as already discussed. In the context of dealing with interference, however, beamforming typically refers to a class of signal processing techniques used for maximizing the signal energy sent to the desired user, while minimize the interference sent towards interfering users. Beamforming may be used to focus energy or may be used to support multiple users through a concept known as spatial division multiple access (SDMA). Beamforming in both forms can be combined with spatial multiplexing to give multiple high data rate streams, although the dimensions used for beamforming do reduce the number of simultaneous data streams that can be transmitted. Typically beamforming for interference reduction is better suited to battling self-cell interference since it requires the complete interference statistics for each user at the transmitter, but in doing so it reduces the other-cell interference. Reducing OCI directly is difficult since it is hard for the base station to acquire appropriate statistical interference knowledge due to the asymmetry between the interference experienced at the mobiles (from other base stations), and the interference experienced at the base station (from other mobiles).

4) Multiuser Diversity: In a data system, when there are $N > 1$ active users, it is possible to take advantage of multiuser diversity by scheduling transmissions to users with good channels [21]. Though originally proposed for single antenna systems, there are several different extensions possible for MIMO systems. The best approaches use a modification of SDMA where multiple users share the same spatial channel (this is sum capacity optimal) though gains are still possible when all the antennas are allocated to a single user. Capacity gains with multiuser diversity grow as $\log \log N$ in Rayleigh fading channels, so the majority of the capacity gains are achieved with just a few users. In wideband channels however, the gain from the SDMA multiuser diversity techniques rapidly decreases due to the abundant frequency diversity. Therefore, multiuser diversity is more effectively exploited in the frequency domain, for example by assigning different subcarriers to different users in an OFDMA system. In summary, multiuser diversity simply provides another form of diversity in the system and can reduce OCI somewhat when combined with transmit power control.

IV. Strategic Approaches to the OCI Problem

As discussed in the prior section and summarized in Table I, the numerous recently proposed methods for mitigating OCI in cellular MIMO systems have some important shortcomings when viewed in a practical context. Although some of these techniques have important merits and are being actively researched and considered, it is useful to consider strategic approaches to handling OCI in cellular systems that do not require real-time information about the other-cell interference. In this section, we overview system-level techniques that can in principle be combined with the previously discussed advanced signal processing algorithms to yield even larger gains, if and when those become viable. A general theme is that network level cooperation can be used to coordinate transmissions resulting in an overall OCI reduction. While these techniques will effectively reduce OCI and increase spectral efficiency for any interference-limited cellular system, the gains are particularly impressive for multicell MIMO systems due to its severe degradation in the face of OCI.

A. Networked MIMO: Base Station Cooperation

We first consider the scenario where the base stations are networked together, and can share different amounts of information with each other. Communication between base stations already occurs in order to coordinate handoffs and other network-level operations.

1) Joint Encoding and "Dirty Paper Coding": If the received interference signals are known to the transmitters, cooperative encoding among neighboring base stations can suppress other-cell interference. Such a joint encoding scheme is an example of so-called dirty paper coding (DPC), which has been shown to achieve the (maximum theoretical) capacity of the multiuser MIMO downlink channel, see e.g. [22]. However, joint encoding is nearly impossible to achieve in practical systems because it requires precise time and phase synchronization of the signals transmitted from multiple base stations, and exact channel knowledge at all the transmitters. Even though cellular systems have base station controllers (BSC) or radio network controllers (RNC), which control multiple base stations, the precise accuracy required for DPC renders this technique as a theoretical upper bound, rather than a practical solution.

2) Cooperative Scheduling: Since instantaneous full channel knowledge is a daunting requirement for a cellular network, a simpler possibility is to allow the base stations to simply take turns transmitting. Recent work, e.g. [23]–[25], has investigated the possibility of neighboring base stations scheduling their transmissions in a cooperative fashion. This is a generalization of the concept of spatial frequency reuse, where simple intercell coordination is used to select appropriate users on a system wide level for each time slot. As a simple example, two base stations should not simultaneously transmit at full power to mobile stations on their mutual cell boundaries.

Just as frequency reuse achieves OCI reduction at the expense of the frequency reuse factor, cooperatively scheduled transmission reduces OCI at the expense of a transmit duty cycle. There are two important advantages of cooperatively scheduled transmission relative to traditional frequency reuse. First, universal frequency reuse (i.e. 1:1 reuse), which simplifies frequency planning. Second, cooperatively scheduled transmission achieves an additional multiuser diversity gain of $\sqrt{\log N}$ if the base stations schedule opportunistically among N neighboring base stations. In contrast to cooperative encoding, cooperative scheduling requires only minimal information (i.e. a single scalar like maximum throughput) to be shared among neighboring base stations, and hence is comparable to handoff in terms of the amount of coordination required between base stations. Additionally, unlike other forms of multiuser diversity that rely on small-scale fading, this macrodiversity gain (from both Rayleigh and lognormal fading) does not compete with other forms of microdiversity (like frequency diversity) and is thus likely to be far more effective in an actual deployment.

For a 2×2 MIMO system with cooperatively scheduled transmission, the capacity gain of cooperatively scheduled transmission over a traditional frequency reuse system ($f = 7$) is shown in Figure 4, where all the users in each cell are assumed to be randomly (uniformly) distributed. As expected, cooperatively scheduled transmission exploits expanded multiuser diversity and achieves higher capacity than traditional frequency reuse ($f = 7$). The expanded multiuser diversity gain in terms of capacity is about 1 bps/Hz. It should be stressed that the main motivation for cooperatively scheduled transmission is not to increase the overall capacity relative to universal time and frequency reuse (the $1/N$ duty factor precludes this), but to increase the effective SINR of the many users that are near cell boundaries. We expect practical base-station cooperation techniques to become popular in the next few years.

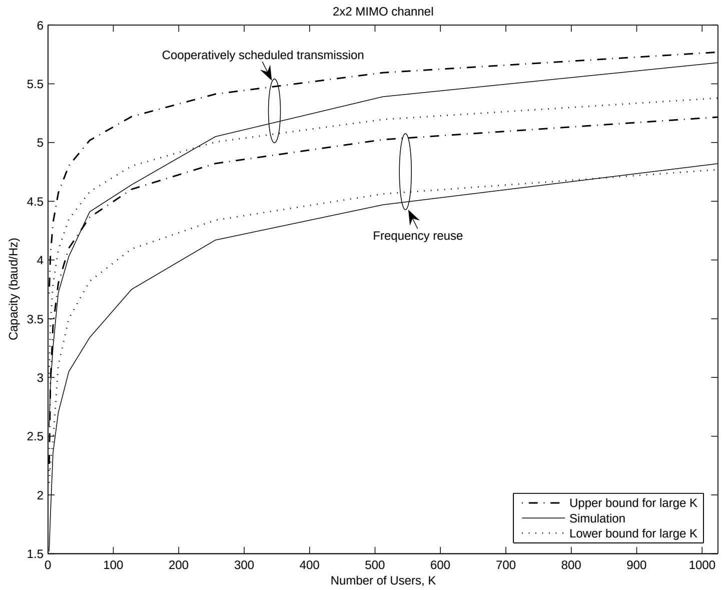

Fig. 4. Shannon capacity of TDMA systems with cooperatively scheduled transmission and with dynamic frequency reuse ($f = 7$). The expanded multiuser diversity gain from cooperatively scheduled transmission is about 1 bps/Hz.

B. Distributed antenna architectures

Another strategic approach to reducing the other-cell interference problem is to adopt a distributed antenna architecture, which has been considered in the past as an effective means of extending coverage, eliminating dead spots, and lowering the blocking probability [26]. In such an architecture, antenna modules are geographically distributed throughout the original cell to reduce the access distance for each node, and are connected to a home base station (or central unit) via dedicated wires, fiber optics, or an exclusive RF link. An example of the distributed antenna cellular structure is given in Figure 5. Distributed antenna systems (DAS) are a low-cost alternative to micro or pico-cells, since typically the distributed antenna units are small (and hence easily mountable) and only require a power amplifier and very minimal other hardware; the vast majority of the processing is performed at the central base station, which has additional advantages in terms of trunking efficiency, maintenance, and ease of

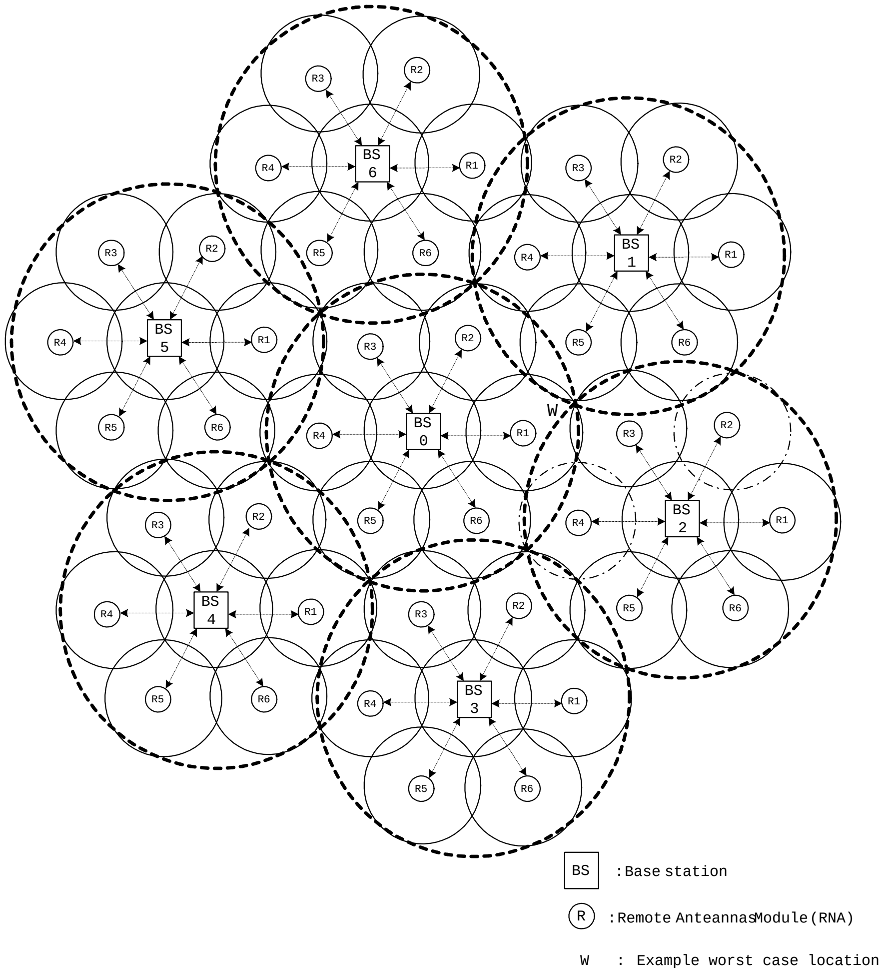

Fig. 5. An example of a distributed antenna cellular architecture for 7 cells, each serviced by 7 total antennas – one at the original base station and 6 remote antenna modules that are spread throughout the cell.

handoff. Numerous equipment manufacturers and service providers are aggressively pursing DAS architectures, which are the subject of a forthcoming book [27].

Since a distributed antenna architecture lowers the aggregate transmit power, it is natural that such a scheme will lower the amount of interference caused to neighboring cells. By transmitting to each user from the antenna closest to it (or more precisely, with the best channel), a form of macroscopic spatial diversity is also introduced to the system, with the net result being a much higher average received SINR, with particularly large gains for users near the distributed antennas

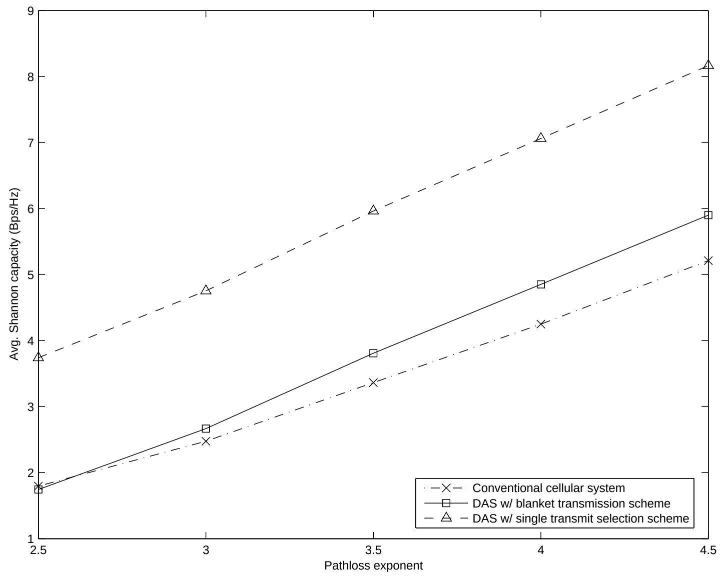

Fig. 6. Average ergodic capacity with CSIR versus the pathloss exponent for one transmit and receive antenna. The capacity gain from distributed antenna systems is dramatic even with universal frequency reuse.

(and hence the cell edges). Figure 6 shows the ergodic capacity of a distributed antenna system with various transmission schemes when the target mobile stations are uniformly distributed and channel state information is available only at the receivers (CSIR). As can be seen, the capacity gain is quite dramatic [28]. This large improvement is due to the fact that three times more users are outside of the radius $1/2R$ within a cell, assuming uniform distribution. Although this was only for a 1×1 system (i.e. $M_r = M_t = 1$), it is intuitive that the capacity gains will be similar for a MIMO system since the transmit powers scale linearly with M_t. The large capacity gain achieved from deploying distributed antennas could be key to the success of large-scale MIMO cellular systems with universal frequency reuse.

V. The Future of MIMO in Wide Area Wireless Networks

This article has attempted to establish that one of the key challenges facing the deployment of MIMO technology in cellular networks is the sensitivity of MIMO receivers to interference. Since cellular systems are inherently interference-limited, this introduces a fundamental conflict. On one hand, system designs should minimize transmit power and maximize spatial reuse in order to reduce the interference caused to neighboring cells. On the other hand, spatial multiplexing MIMO systems by nature increase the number of interfering sources, and are effective mostly in high SINR.

Traditional approaches to multicell interference management have included static techniques frequency reuse, sectoring, and spread spectrum. Frequency reuse and spread spectrum are not bandwidth efficient solutions, which defeats the purpose of MIMO, while sectoring is effective but is already near its limits. Recent research approaches to this difficult problem have focused on advanced signal processing techniques at the receiver and transmitter as a means of reducing or cancelling the perceived interference. As this article has documented, however, most of these techniques suffer from important practical shortcomings in terms of complexity and required channel information that make their successful application to cellular systems unlikely in the near to medium term.

As an alternative, this paper has advocated strategic approaches that require very little channel knowledge and effectively reduce other-cell interference through macro-diversity. Allowing for some simple back-channel communication among neighboring base stations, as is typically the case due to the need to coordinate handoffs and other operations, base station cooperation will be essential for interference reduction. Even larger gains in interference reduction are possible with a distributed antenna architecture, which are already being widely considered as a means to extend coverage. In the future, we expect that a suite of several of the techniques overviewed in this paper are likely to be deployed in order to manage the interference problem in spatial multiplexing wireless networks.

VI. ACKNOWLEDGEMENTS

The authors gratefully acknowledge feedback and input from M. Shafi (NZ Telecom), A. Ghosh (AT&T Labs), E. Onggosanusi (TI), S. Yi (SOLiD), H. Dai (N.C. State), N. Jindal (Minnesota), and S. Talwar (Intel). We would also like to thank the anonymous reviewers for their helpful suggestions.

TABLE I
SUMMARY OF POSSIBLE OTHER-CELL INTERFERENCE (OCI) MITIGATION TECHNIQUES

Technique	Benefit	Key shortcomings	Prospects
Frequency reuse	Reduce OCI very simply and effectively	Low spectral efficiency, frequency planning	Not promising as a long-term solution
Maximum Likelihood MUD	Optimum co-reception of signal and interference	Very high complexity, OCI-awareness	Moore's law will help complexity, but prohibitive in near future
MMSE MUD	Suppresses OCI with much lower complexity than ML	Requires awareness of OCI, many mobile antennas. Simpler versions have only modest performance gain.	Requires instantaneous OCI knowledge, under present investigation by industry
OCI-blind MMSE	Like ZF spatial receiver with lower noise enhancement	Enhances OCI rather than suppressing it: very poor performance	Will provide only incremental gain (10-25%), but likely to be implemented
Other-cell interference cancellation	Good performance vs. complexity	Complexity still high, awareness and accuracy of OCI knowledge crucial	Promising in long- to provide additional gain over OCI-blind receivers
Stream Control	Reduces OCI; increases robustness	Lowers the data rate	Adaptive stream control is feasible and useful.
Multiuser Diversity	Decreases required transmit power or increases data rate	Competes with other forms of diversity like frequency diversity. log log N growth – i.e. rapidly diminishing returns.	Like to be useful in terms of scheduling, but not very effective for OCI reduction
Cooperative Encoding, i.e. Dirty Paper Coding	Optimal performance in theory	Requires very accurate channel knowledge and real-time inter-cell coordination	Unlikely to be practical in foreseeable future, if ever.
Closed-loop MIMO diversity	Achieves optimum diversity performance	Sacrifices spatial dimensions for multiplexing, channels known at Tx	Likely to be implemented, can lower OCI somewhat.
Beamforming	Reduces OCI	Sacrifices spatial dimensions, channels known at Tx	Has important merits, but implementation difficulties
Cooperative Transmission	Reduce OCI, multiuser diversity gain relative to frequency reuse	Requires simple cooperation between base stations	Feasible in the short-term
Distributed Antenna Systems	Reduce OCI through lowered transmit power; better coverage; ease of maintenance	Requires new infrastructure deployment paradigm	Feasible in the short-term with large infrastructure investment

References

[1] S. M. Alamouti, “A simple transmit diversity technique for wireless communications,” *IEEE Journal on Sel. Areas in Communications*, vol. 16, no. 8, pp. 1451–1458, Oct. 1998.

[2] V. Tarokh, N. Seshadri, and A. Calderbank, “Space-time codes for high data rate wireless communications: Performance criterion and code construction,” *IEEE Trans. on Info. Theory*, vol. 44, no. 2, pp. 744–765, Mar. 1998.

[3] R. Derryberry, S. Gray, D. Ionescu, G. Mandyam, and B. Raghothaman, “Transmit diversity in 3G CDMA systems,” *IEEE Communications Magazine*, vol. 40, no. 4, pp. 68–75, Apr. 2002.

[4] E. Telatar, “Capacity of multi-antenna Gaussian channels,” *European Trans. Telecommun.*, vol. 6, pp. 585–595, Nov.-Dec. 1999.

[5] R. W. Heath, S. Sandhu, and A. Paulraj, “Antenna selection for spatial multiplexing systems with linear receivers,” *IEEE Communications Letters*, vol. 5, no. 4, pp. 142–44, Apr. 2001.

[6] A. F. Molisch and M. Z. Win, “MIMO systems with antenna selection,” *IEEE Microwave Magazine*, vol. 5, no. 1, pp. 46–56, Mar. 2004.

[7] D. J. Love, R. W. Heath Jr., W. Santipach, and M. L. Honig, “What is the value of limited feedback for MIMO channels?” *IEEE Communications Magazine*, vol. 42, no. 10, pp. 54–59, Oct. 2004.

[8] D. Gesbert, M. Shafi, D. Shiu, P. J. Smith, and A. Naguib, “From theory to practice: An overview of MIMO spacetime coded wireless systems,” *IEEE Journal on Sel. Areas in Communications*, vol. 21, no. 3, pp. 281–302, Apr. 2003.

[9] R. Blum, “MIMO capacity with interference,” *IEEE Journal on Sel. Areas in Communications*, vol. 21, no. 5, pp. 793–801, June 2003.

[10] S. Catreux, P. F. Driessen, and L. J. Greenstein, “Attainable throughput of an interference-limited multiple-input multiple-output (MIMO) cellular system,” *IEEE Trans. on Communications*, vol. 49, no. 8, pp. 1307–1311, Aug. 2001.

[11] A. Forenza, M. R. McKay, I. B. Collings, and R. W. Heath, “Switching between OSTBC and spatial multiplexing with linear receivers in spatially correlated MIMO channels,” in *Proc., IEEE Veh. Technology Conf.*, Melbourne, Australia, May 2006.

[12] R. Valenzuela, “Fundamental limits and evolution of broadband wireless access networks,” Texas Wireless Symposium, Austin, TX, available at http://www.ece.utexas.edu/wncg/events/symposium2005/, Nov. 2005.

[13] W. Choi and J. G. Andrews, “Spatial multiplexing in cellular MIMO-CDMA systems with linear receivers: Outage probability and capacity,” *IEEE Trans. on Wireless Communications*, to appear, available at: www.ece.utexas.edu/∼jandrews.

[14] H. Dai, A. Molisch, and H. Poor, “Downlink capacity of interference-limited MIMO systems with joint detection,” *IEEE Trans. on Wireless Communications*, vol. 3, no. 2, pp. 442–453, Mar. 2004.

[15] A. Burg, et al, “Performance tradeoffs in the VLSI implementation of the sphere decoding algorithm,” in *IEE 3G Mobile Communication Technologies Conference*, London, UK, Oct. 2004.

[16] W. Choi, J. G. Andrews, and R. W. Heath, “Multiuser antenna partitioning for MIMO-CDMA systems,” *IEEE Trans. on Veh. Technology*, to appear, available at: www.ece.utexas.edu/∼jandrews.

[17] J. G. Andrews, “Inteference cancellation for cellular systems: A contemporary overview,” *IEEE Wireless Communications Magazine*, vol. 12, no. 2, pp. 19–29, Apr. 2005.

[18] A. Scaglione, P. Stoica, S. Barbarossa, G. B. Giannakis, and H. Sampath, “Optimal designs for space-time linear precoders and decoders,” *IEEE Trans. on Signal Processing*, vol. 50, no. 5, pp. 1051–1064, May 2002.

[19] R. Blum and J. Winters, “On optimum MIMO with antenna selection,” *IEEE Commun. Lett.*, vol. 6, no. 8, pp. 322–324, August 2002.

[20] R. W. Heath and D. J. Love, “Multimode antenna selection for spatial multiplexing systems with linear receivers,” *IEEE Trans. on Signal Processing*, vol. 53, pp. 3042–3056, August 2005.

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

[22] A. Goldsmith, S. Jafar, N. Jindal, and S. Vishwanath, “Capacity limits of MIMO channels,” *IEEE Journal on Sel. Areas in Communications*, vol. 21, no. 5, pp. 684–702, June 2003.

[23] W. Choi and J. G. Andrews, “The capacity gain from base station cooperative scheduling in a MIMO DPC cellular system,” in *Proc., IEEE Intl. Symposium on Information Theory*, Seattle, WA, June 2006.

[24] H. Zhang and H. Dai, “Co-channel interference mitigation and cooperative processing in downlink multicell multiuser MIMO networks,” *European Journal on Wireless Communications and Networking*, pp. 222–235, 4th Quarter 2004.

[25] G. J. Foschini, H. Huang, K. Karakayali, R. A. Valenzuela, and S. Venkatesan, “The value of coherent base station coordination,” in *Proc., Conference on Information Sciences and Systems (CISS)*, Johns Hopkins University, Mar. 2005.

[26] A. Saleh, A. J. Rustako, and R. S. Roman, “Distributed antennas for indoor radio communications,” *IEEE Trans. on Communications*, vol. 35, pp. 1245–51, Dec. 1987.

[27] H. Hu, Y. Zhang, and J. L. (Editors), *Distributed Antenna Systems: Open Architecture for Future Wireless Communications*. CRC Press, To appear, 2007.

[28] W. Choi and J. G. Andrews, “Downlink performance and capacity of distributed antenna systems in a multicell environment,” *IEEE Trans. on Wireless Communications*, Jan. 2007.