Dealing with Loud Neighbors: The Benefits and Tradeoffs of Adaptive Femtocell Access

David Choi, Pooya Monajemi, Shinjae Kang, and John Villasenor
Electrical Engineering Dept., University of California, Los Angeles
dschoi@ee.ucla.edu, pmonajemi@ucla.edu, shinjae@ee.ucla.edu, and villa@icsl.ucla.edu

Abstract—**Femtocells are low-power, very small-service-area (e.g. home or office environment) cellular base stations that will significantly impact the cellular landscape in the next several years. One of the most important open technical issues related to femtocells concerns the impact on system performance of different policies regarding who is allowed to connect to a femtocell. In the present paper, interaction between mobile stations (MS) that are near to, but not necessarily communicating with, femtocells is explored. It is shown that an adaptive femtocell access policy that takes specific account of the instantaneous loads on the network can lead to improved performance over a completely open, or completely closed approach.**

I. INTRODUCTION

Femtocells, which are also sometimes referred to using the term "Home NodeB" (HNB), are low-power, low-cost cellular base stations designed to serve a very small area, such as a home or office environment. There are a number of economic factors that argue strongly in favor of femtocells, including the capacity improvements that are enabled by allowing mobile stations (MS) to transmit at very low power to a femtocell that in turn utilizes a high speed internet connection as the backhaul. There is high interest in femtocells among carriers, equipment manufacturers, and other participants in the cellular industry, and many analysts expect that femtocells will significantly impact the cellular landscape over the next several years.

One of the most interesting open technical questions concerns how best to optimize the performance of networks including femtocells in light of the interactions between femtocells and MSs that may be near, but not communicating with, the femtocell. Like Wi-Fi access points, femtocells can be operated in either in a closed subscriber group (CSG) manner (also called "closed access"), in which only MSs with proper authorization can connect, or under an open access policy, in which any MS can connect. If an MS is very near to a femtocell but unable to connect to it, it will need to connect to a regular ("macro") base station that may be hundreds of meters away. This creates the potential for a "loud neighbor" effect, in which the strong signal from this MS contributes significantly to the noise environment in the neighborhood of the femtocell.

One way to circumvent this problem is to use a separate, dedicated channel for the femtocells at a different frequency from that employed by the macro base stations in a network. However, in many cases this will not be possible because a single carrier may not own enough spectrum. For example, in many regions carriers have licensed three 5-MHz frequency bands, and in order to maximize revenue, will fully utilize this spectrum for traditional macro cellular services. To the extent that femtocells are added in such in environment, they will have to operate in a co-channel manner, sharing the same spectrum as the macrocells and thereby creating a potential for the interference effects described above. In dense urban areas, this effect can be significant. For example, in office complexes and apartment buildings in an urban core, there can easily be several hundred people within a few tens of meters of a given location. Even after adjusting for multiple carriers, frequency bands, and the expected use patterns of individual MS users, there can easily be significant numbers of active MSs near to, and at the same frequency as, a given femtocell.

The issue of how best to manage co-channel femtocells has been given significant attention in the 3GPP industry group, with particular attention to the performance under the two scenarios of open or closed access. These studies have explored issues including downlink capacity for networks with open vs. closed access [1], the effect of the number femtocells on the capacity of a macrocell [2], and examinations of interference impacts involving various combinations of femtocells and macrocells [3]. The conclusion from these studies is that a closed co-channel femtocell can lead to significant interference problems for all parties, particularly if the femtocell does not adaptively change its transmit power in order to minimize its interference on existing networks [4].

The goal of this paper is to explore some of the subtleties involved in open access, and in particular to examine the tradeoffs associated with different levels of open access. The underlying premise of this paper is that the value of a femtocell both to its owner and to the overall network will be maximized if the level of open access is adaptively controlled as a function of factors including the instantaneous load on the femtocell. While a closed femtocell can be problematic as noted above, in an environment with a high density of MSs a completely open femtocell can also suffer problems, because it will potentially force the sharing of limited femtocell wireless bandwidth and internet backhaul capacity among significant numbers of MSs,

This work was supported in part by the Office of Naval Research (Contract number N00014-06-1-0253) and in part by the National Science Foundation (Grant number CCF-0541453).

978-1-4244-2324-8/08/$25.00 © 2008 IEEE.

with obvious detrimental consequences to the service level available to the owner of the femtocell.

The remainder of this paper is organized as follows. In Section II we present an analytical framework that addresses the capacity of femtocells embedded in a macro network. In Section III, we describe a simulation environment that models both legacy WCDMA as well as high speed downlink packet access (HSDPA), taking into account the specific uplink and downlink messages that occur in each 2 msec time slot. Section IV shows simulation results with respect to open vs. closed femtocell policies. Conclusions are offered in Section V.

II. CAPACITY CALCULATIONS IN A NETWORK OF MACROCELLS AND FEMTOCELLS

Femtocells will be deployed in environments that are inherently complex, involving a mix of circuit-switched traffic from legacy WCDMA systems and packet-based high speed access (HSDPA) communications. In the interest of providing an analytical foundation for a network comprising macrocells and femtocells, an approach for evaluating the downlink capacity is presented. In [5], Hiltunen *et al.* provide a method for evaluating the downlink capacity of a cellular system containing only legacy WCDMA users. In the approach presented in [6] a similar method is presented to evaluate HSDPA capacity in a system comprised of both legacy and HSDPA users, studied separately for femtocells and macrocells. Our approach considers scenarios where both femtocells and macrocells are present, and can be used to study the tradeoffs that exist in such situations.

Since the high speed downlink data rates assigned to scheduled users are a monotonically increasing function of the carrier to interference ratios achievable at the location of the mobile station, the basis of the approach is to evaluate the distribution of these carrier to interference ratios among the locations that are supported by the base station under consideration. The ratio is computed at a given location x (where x is a vector) with regard to an assigned base station j as follows:

$$\gamma_{x,j} = \frac{\frac{P_{tot,j} \cdot \beta_{HS}}{L_{x,j}}}{\frac{\alpha \cdot P_{tot,j}}{L_{x,j}} + \sum_{i=1, i \neq j}^{B} \frac{P_{tot,i}}{L_{x,i}} + N} \qquad (1)$$

where $P_{tot,j}$ is the total power from base station j, β_{HS} is the fraction of the base station power assigned to HSDPA, $L_{x,j}$ is the link loss from location x to base station j, α is the non-orthogonality factor, B is the total number of base stations considered, and N is the noise power at the mobile station. The underlying assumption in the above relation is that base stations utilize their total available power. An estimate of the HSDPA throughput for a cell is as follows:

$$T_j = \frac{\int_{x \epsilon R_j} f(\gamma_{x,j}) \cdot p(\gamma_{x,j}) \cdot \mu(x) dx}{\int_{x \epsilon R_j} \mu(x) dx} \qquad (2)$$

where T_j is the throughput for base station j, R_j is the region supported by base station j, $f()$ is a mapping function between the carrier to interference ratios and the assigned HSDPA data rates, $p()$ is a factor determined by the scheduling algorithm (constant for fair scheduling) and $\mu()$ is an indicator of user density in a location.

For validation purposes we have used the above analytical framework to examine the network capacity under a range of different scenarios in which both the access policy and the locations of the macro and femto base stations are varied. Distributions for γ were obtained via numerical integration over a large number of realizations. In particular, we have considered a spectrum of access policies in which outdoor MSs are assigned a fixed probability of being able to connect to femtocells (0% corresponds to a completely closed system and 100% corresponds to an open system). The obtained γ distributions for different access policies show that open access policies in general provide better overall C/I performance in the network, especially for the cell edge users. As discussed further below, the results are consistent with those described in the next section with respect to variations in the HSDPA data capacity as a function of different user density levels and the femtocell access policy.

III. SIMULATOR DESCRIPTION

The analytical framework above provides an important "sanity check" on general capacity characteristics. However, for examining scenarios involving a given set of users at a specific set of coordinate locations it is necessary to utilize a comprehensive network simulation environment. Such modeling is an essential tool in understanding the behavior of complex cellular environments including macro and femtocells, and supporting a mix of circuit-switched legacy voice traffic as well as packet-switched HSDPA data traffic. In particular, proper modeling of the HSDPA channels requires a dynamic simulator in which the network behavior is evolved over a series of time slots. The simulation environment utilized in the work performed here was constructed by starting with a static simulator for legacy WCDMA circuit-switched voice communications developed by Laiho and Wacker [7] and building on top of it a dynamic simulator to enable handling of packet data. These changes included 1) the addition of mechanisms to handle the high speed downlink control, uplink control, and downlink data channels, and 2) the addition of the time slotted structure and scheduling methods to introduce dynamism.

A. High Speed Downlink Control Channel

The high speed downlink control channel (HS-SCCH) uses a downlink data rate of 18.5 kbps in accordance with [8]. The E_b/N_0 requirement of this channel is set to be 12 dB. The computed power in the base station is compared to the link power limitation (as is done in the legacy channel model) to check for base station link outages. It is also summed with legacy powers to find the total base station power. This quantity is compared with the maximum total base station power limitation to identify potential outages and is subsequently used in the computation of interference levels.

B. High Speed Uplink Control Channel

The high speed uplink control channel (HS-DPCCH) is used to transmit acknowledgements (ACK) and channel quality indicator (CQI) information during the first 0.67 msec and remaining 1.33 msec of each 2 msec subframe respectively. The uplink data rate is 15 kbps for both the ACK and the CQI. The E_b/N_0 requirement of the channel during ACK is estimated from [9] to be 6.9 dB for a 1% error rate, and the E_b/N_0 requirement of the channel during CQI is set to 3.9 dB. The computed power in the mobile station is added to the legacy channel powers and compared to the maximum mobile station transmit power limitation.

C. Modeling the High Speed Downlink Data Channel

The high speed downlink data channel (HS-PDSCH) differs from conventional WCDMA channels in that it utilizes adaptive modulation and coding (AMC) as opposed to the power control mechanisms Instead of computing an optimum power level, the base station maximizes the transmission data rate for a given HS-PDSCH channel power level by selecting an appropriate modulation scheme (QPSK vs. 16-QAM), coding rate, and number of multicode channels utilized (out of 15 total) based on the available power, link loss data, and interference at the receivers.

There are many possible design choices regarding how many high speed downlink data channels are used at one time. These choices involve a large number of tradeoffs, some obvious, and others much less so, that impact the overall effectiveness of the high speed data delivery. In [10], it has been proposed that the optimal number of mobile stations per time interval is one, and that is the approach that has been used in the experimental results presented here, though this is obviously an area of interest for further exploration.

The HS-PDSCH channel power is modeled as a power assignment in the base station. This power level is chosen as the lower of the available base station power and 30% of the total base station power. Adaptive Modulation and Coding (AMC) is then used to calculate the optimal high speed downlink bit rate.

D. Time Slotted Simulator Implementation for HSDPA

The high speed downlink control, uplink control, and downlink data channels are active during different stages of high speed transmission. High speed downlink control information, which is necessary for signaling the transport format and the multicode information used for decoding high speed downlink data, is sent first. Downlink data then arrives through the high speed downlink data channel. After data has been received and processed, the mobile station sends ACK and CQI messages that are then used by the base station for the next transmission.

Time is considered in 667 μsec slots and each of the three channels is either included or excluded from calculations in each time slot. Each time slot is independently simulated using the augmented static simulator to identify the optimum power levels at each point in time. The downlink control signals are sent in the first 3 slots, the high speed downlink packet is sent in 9 slots, the uplink ACK utilizes 1 slot, and the uplink CQI data spans 2 slots (Fig. 1). The time between high speed downlink control and high speed downlink data is set to be 3 slots to allow the mobile station to process the downlink control data, and the time between receiving the high speed downlink data and sending the high speed uplink control is set to be 6 slots to allow for propagation delay and processing on the mobile station.

A scheduler performs the task of choosing the high speed users at each base station for each packet, and determines the channels that are active for each slot. To reflect the assignment of multicodes to exactly one mobile station at a time, only one mobile station can be the high speed user per base station. For the simulation, a round robin scheduler is used wherein all mobile stations connected to a given base station that desire high speed downlink data take turns being the high speed user. In order to avoid synchronous transmission of ACK signals in all base stations, a random offset of between 0 and 8 time slots is introduced at each base station.

All users are assumed to have voice channels active at all times and to be always ready for receiving high speed data. They are also assumed to be capable of receiving and processing all available AMC combinations and any number of multicodes (up to 15).

For the results described here, the simulation duration spans 3000 time slots, which is equivalent to 2 seconds. During this period mobile stations are assumed to have constant link loss. A given high speed packet transmission is classified as a failure when the mobile station is put to outage in any of the time slots except CQI.

IV. RESULTS

There is clearly an enormous range of scenarios that could be considered using the modeling environment described above. In the interest of clarity and brevity, we limit consideration here to one specific scenario that is very effective at illustrating tradeoffs in femtocell access policy. We have considered a scenario in which we define a number of randomly distributed “indoor” and “outdoor” environments with macrocells, femtocells, mobile stations. An “indoor” environment is a circular region 20 m in diameter surrounded by walls that attenuate all signals traversing them by 15 dB. Inside each indoor environment are two indoor users. All other locations on the map are “outdoor” environments. This is shown in Fig. 2, which contains 32 “indoor” environments designated by circles (indoor environments that include a femtocell) and an additional 32 indoor environments designated by squares (indoor environments with no femtocell). There are also 21 macrocell base stations, designated by lines, each serving a single 120-degree sector. Finally, there are 456 mobile stations designated by dots. The two users inside each femtocell are within the closed subscriber group of that femtocell, and are referred to as the “CSG MSs”. In addition, there are clusters of eight MSs that are near to but outside the walls of the indoor environments containing the femtocells, and for which the common pilot channel (CPICH) signal received from the

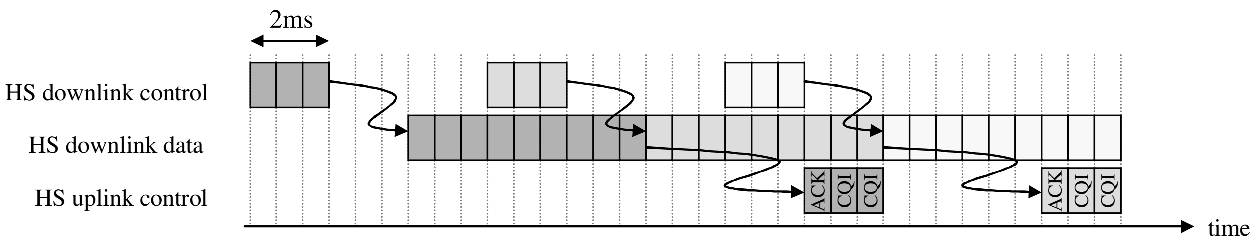

Fig. 1. High speed downlink control, high speed downlink data, and high speed uplink control channel timing for the time slotted simulator implementation.

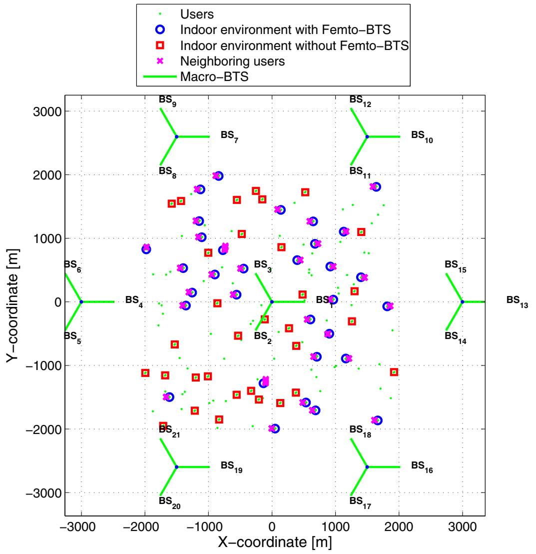

Fig. 2. Positions of MSs, indoor environments, femtocell locations, and macro base stations.

femtocell is stronger than that received from the macro base station. MSs that fall into this category are referred to as "neighbors", as they could potentially connect to the femtocell under an open access scenario.

The service results are shown in Fig. 3. Different positions on the horizontal axis represent different policies for femtocell access. At the left end of the axis, the case with no femtocells is considered. In this case all of the MSs, including the two MSs that will later be part of the CSG, connect (or attempt to connect) to the macro base stations. The next position on the horizontal axis is labeled 2CSG, and considers that performance when the femtocells are present and the access is closed so that only the two CSG mobiles can utilize them. As one moves further to the right on the axis, successively larger numbers of neighbors are granted access to the femtocells.

The vertical axis gives the number of MSs associated with the macro and femtocell that are served and in outage. For example, when there are no femtocells, 406 of the 456 MSs receive service using the macro base station. When the femtocells are present and limited to the fully closed ("2CSG") approach, 212 MSs are served by the macro base stations and 177 MSs are not served, while 67 MSs are served by the femtocells. Thus, it can be seen that in this scenario, introducing femtocells with a closed access policy results in a decrease in service. This is due to the interference that the femtocells and the femtocell users are generating, which is sufficiently harmful to the macrocell users to cause outages. However, as the femtocell is opened up to allow more users, the overall service improves. It can be seen that allowing 6 neighbors into the femtocells results in similar service to the macrocell-only scenario, and that opening up access policy further results in even better service.

The throughput results are shown in Fig. 4. Just as in Fig. 3, different positions on the horizontal access represent different policies for femtocell access. The vertical axis gives the average throughput for those MSs connecting to the macrocell (dark bars), the femtocell (light bars), and the average over all mobiles in the system (solid line with triangles). Notably, in the fully closed approach, the average femtocell MS throughput is maximum at 3.17 Mbps. As the access policy is changed toward a more open policy, the average femtocell MS throughput decreases. This is because the femtocell needs to share its bandwidth among the neighbors that receive access to the femtocell. Moreover, as more users are allowed into the femtocell, the number of users served per each macrocell decreases, so the throughput enjoyed by the macrocell users increases.

The optimal operating point for this scenario is dependent on perspective. From the standpoint of the femtocell owner, a closed access policy is optimal, because all CSG MSs are successfully able to connect and their throughput is maximized. From the standpoint of those users in the system who are not using the femtocell, a fully open approach (in which all eight neighbors are granted access) is optimal, because in this case the femtocell is providing maximal relief to the load on the macrocell, thereby enabling those MSs still using the macro base station to receive higher throughput. A limited access policy of allowing 6 neighbors into the femtocell results in similar service levels to the macrocell-only scenario, but with a significantly improved throughput.

These outage results are consistent with Monte Carlo simulations based on the analytical approach described in section II above, which confirm that from an analytical standpoint, the outage probabilities can be significantly reduced with progressively more open femtocell access policies.

In sum, allowing access to six neighbors maintains the

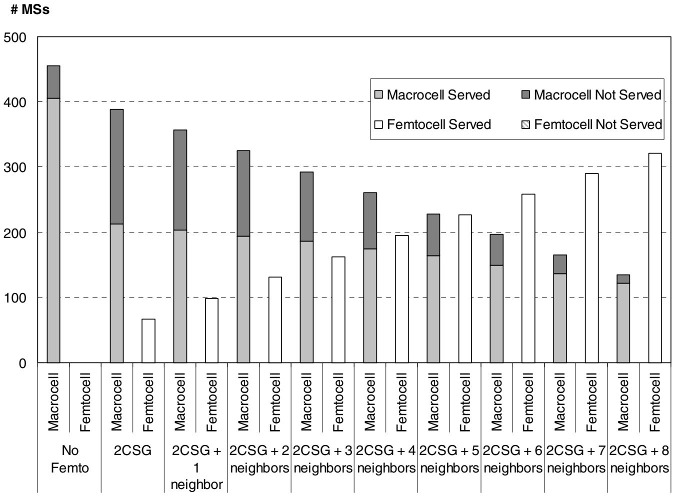

Fig. 3. Service as a function of different femtocell access policies, ranging from no femtocell, a fully closed femtocell (“2CSG” label), and successively more open access as one moves to the right on the horizontal axis.

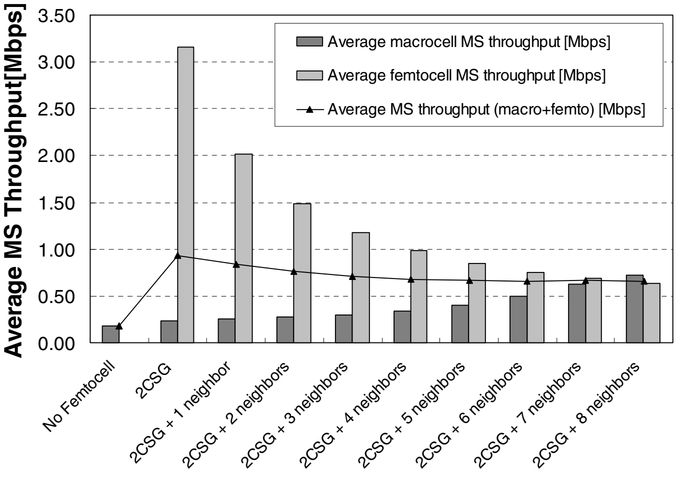

Fig. 4. Average throughput as a function of different femtocell access policies, ranging from no femtocell (e.g. the femtocell is not present or is powered down) at the left end of the axis, a fully closed femtocell (“2CSG” label), in which only the the two MSs that are members of the CSG are allowed to connect), and successively more open access as one moves to the right on the horizontal axis.

same level of outage performance while improving throughput performance for all the MSs relative to a system with no femtocell or a fully closed femtocell. In the particular scenario explored in Fig. 2, the presence of the femtocell is a win-win - the femtocell owners benefit, as do the other MSs, even those not utilizing the femtocell. In other simulations, we have found that the choice of optimum operating point depends very strongly on the specific scenario, and in addition on the metric chosen for defining “optimum”. For example, additional scenarios have shown benefits of a limited access policy at different optimal operating points, allowing 5, 6, or 7 neighboring users. More important however than the specific results for any one scenario is the general result that the performance of different access policies is highly dependent on the specifics of the scenario, and that all participants in the system have a strong incentive to make sure that access strategy is well managed.

V. CONCLUSIONS

This paper has examined tradeoffs relating to access policies for femtocells that share frequencies with the macrocells in which they are embedded. These issues have been considered both using a generalized analytical capacity description as well as through development of a network modeling environment specifically designed to allow consideration of femtocells. The specific simulation scenario described here, as well as others we have examined, show that the impact of specific access decisions is highly dependent on the specific configuration of the femtocells, nearby macro base stations, and MSs. Completely closed access can disrupt service because of the interference generated by MSs that are near to, but prohibited from attempting to connect to, a femtocell. Completely open access can also be problematic due to the loss in data rate arising from sharing limited femtocell bandwidth among potentially large numbers of users. Rather than completely or closing or opening access, better performance can be achieved by opening the femtocell to the best operating point for a given set of goals concerning connectivity and throughput.

REFERENCES

[1] Nortel, Vodafone, “Open and closed access for home NodeBs,” 3GPP, Athens, Greece, 3GPP document reference R4-071231, Aug. 2007, 3GPP TSG-RAN WG4 Meeting 44.
[2] Nokia Siemens Networks, “Initial home NodeB coexistence simulation results,” 3GPP, Orlando, USA, 3GPP document reference R4-070902, Jun. 2007, 3GPP TSG-RAN WG4 Meeting 43bis.
[3] Orange, “Home BTS consideration and deployment scenarios for UMTS,” 3GPP, Kobe, Japan, 3GPP document reference R4-070825, May 2007, 3GPP TSG RAN WG4 Meeting 43.
[4] “3GPP TR25.8200 v1.0.0 release 8 (11-2007).”
[5] K. S. Gilhousen, I. M. Jacobs, R. Padovani, A. J. Viterbi, L. A. Weaver, Jr., and C. E. Wheatley, III, “On the capacity of a cellular CDMA system,” *IEEE Trans. on Vehicular Technology*, vol. 40, no. 2, pp. 303–312, May 1991.
[6] K. Hiltunen, B. Olin, and M. Lundevall, “Using dedicated in-building systems to improve HSDPA indoor coverage and capacity,” in *Proc. of IEEE VTC 2005-Spring*, vol. 4, May 2005, pp. 2379–2383.
[7] J. Laiho and A. Wacker, “Radio network planning process and methods for WCDMA,” *Annals of Telecommunications*, vol. 56, no. 5-6, pp. 317–331, 2001.
[8] “3GPP TS 25.212 v7.5.0 (2007-05).”
[9] “3GPP TS 25.104 v7.7.0 (2007-06).”
[10] A. Bedekar, S. Borst, K. Ramanan, P. Whiting, and E. Yeh, “Downlink scheduling in CDMA data networks,” in *Proc. of IEEE Global Telecommunications Conference*, vol. 5, Dec. 1999, pp. 2653–2657.