A Simple Greedy Scheme for Multicell Capacity Maximization

Saad G. Kiani*, David Gesbert*, Jan Egil Kirkebø†, Anders Gjendemsjø‡, Geir E. Øien‡
*Mobile Communications Department, Eurecom Institute,
06560 Sophia Antipolis, France. E-mail: {kiani, gesbert}@eurecom.fr
†Department of Informatics, University of Oslo,
P.O. Box 1080, N-0316 Oslo, Norway. E-mail: janki@ifi.uio.no
‡Dept. of Electronics and Telecom., Norwegian Univ. of Science and Technology,
7491 Trondheim, Norway. E-mail: {gjendems, oien}@iet.ntnu.no

Abstract— **We study joint optimization of transmit power and scheduling in a multicell wireless network. Despite promising significant gains, this problem is known to be NP-hard and thus difficult to tackle in practice. However, we show that this problem lends itself to analysis for large wireless networks which allows simpler modeling of inter-cell interference. We introduce a low complexity greedy algorithm that is efficient for large networks. As the number of users per cell increases, the solution converges to all cells being active and employing maximum SINR scheduling, which can be implemented in a distributed manner. Using simulation parameters equivalent to those used in realistic wireless networks we show that the scheme, though simple, exhibits substantial gains over existing resource allocation schemes.**

I. Introduction

Full reuse of spectrum, in any of the dimensions allowed by multiple access schemes (time or frequency slots, codes etc.), is envisioned to offer much greater capacity in wireless data networks. System level performance, however, is adversely affected by an increased, sometimes unbearable, level of interference due to aggressive reuse of the spectral resource. Inter-cell interference mitigation through smart power allocation will therefore play an important role in realizing the benefits of full spectrum reuse. Similarly, recently developed dynamic resource management techniques exploiting multi-user diversity [1] will also offer increased system capacity. Clearly multi-user diversity gain is achieved at the expense of throughput fairness, which may be restored by modifying the scheduling criteria in one of several possible manners [2]. It is evident that allocating power and scheduling users jointly is the key to achieving the maximum gain.

A fundamental point in the joint multicell power allocation and scheduling problem is that a huge number of degrees of freedom (governed by the number of users times the number of possible power allocations) can be potentially used to maximize the *network capacity* in an interference-limited setting. However, despite the potential of significant gains, this problem is NP-hard. An optimal solution will involve finding the power vector and set of scheduled users that will maximize network capacity.

A number of approaches to this problem can be envisioned which may be either centralized or distributed. Attempting to find the optimal solution, e.g. through an exhaustive search, requires centralized processing with knowledge of the complete system parameters. This is not a trivial task given the enormous size of the optimization space. Alternatively, by reformulating the objective function, convex optimization techniques may be applied [3], though this is still centralized and not optimal with respect to system capacity. A completely distributed approach can also be considered [4] which requires knowledge of local channel information. The point is that one can approach this problem in many different ways with varying degrees of complexity to obtain different amounts of capacity gain.

In this paper we use the *interference-ideal* network model [5] and recent results [6] to propose a greedy algorithm for power allocation and user scheduling with the goal of maximizing the network capacity. The algorithm is based on the idea of switching off some cells to decrease system interference and thus obtain a system-wide gain. Simulation results show that as the number of users increases, the solution converges to transmitting with full power in all cells, and scheduling the user with best SINR in each cell. The advantage of this result is that it provides a distributed solution to the joint power allocation and scheduling problem for networks with many users. The proposed scheme is simple in nature, and by using simulation parameters equivalent to those in realistic wireless networks our scheme exhibits substantial gains over existing resource allocation schemes.

In Section II we outline the system model. We then proceed to formulate the joint power allocation and scheduling problem in Section III, and give the objective function we wish to optimize. In Section IV previous results are extended and a greedy algorithm for power allocation and user scheduling is proposed. The paper is concluded in Section V with numerical results comparing the performance of the greedy algorithm with traditional reuse and full reuse schemes.

II. System Model

We consider a multicell system in which many access points (APs) communicate with user terminals (UTs) over a coverage area. We are particularly interested in the downlink in which the APs send data to the UTs. We assume that the system employs the *same spectral resource* in each cell, giving rise to an interference-limited system (Fig. 1), and power control is used in the network in an effort to preserve power and limit interference and fading effects. This results in a peak power constraint $P_{\max}$ at each AP. Within each cell, we consider a multiple access scheme in which an orthogonally divided resource (e.g. codes, time, frequency etc.) is used to separate the transmissions to the cell users. Each cell user is allocated an intra-cell orthogonal resource slot, but due to full reuse the user "sees" interference from all neighboring co-channel cells.

A. Signal Model

The downlink of the multicell system above is considered, with N cells and U_n users randomly distributed over each cell n. To simplify notation we focus on the single antenna case. Denoting the channel gain between any arbitrary AP i and user u_n in cell n by $G_{u_n,i} \in \mathbb{R}^+$, and assuming perfect channel state information (CSI) at the receiver, the received signal Y_{u_n} at the user is given by

$$Y_{u_n} = \sqrt{G_{u_n,n}} X_{u_n} + \sum_{i \neq n}^{N} \sqrt{G_{u_n,i}} X_{u_i} + Z_{u_n}$$

where X_{u_n} is the signal from the serving AP, $\sum_{i \neq n}^{N} \sqrt{G_{u_n,i}} X_{u_i}$ is the sum of interfering signals from other cells, and Z_{u_n} is additive white Gaussian noise.

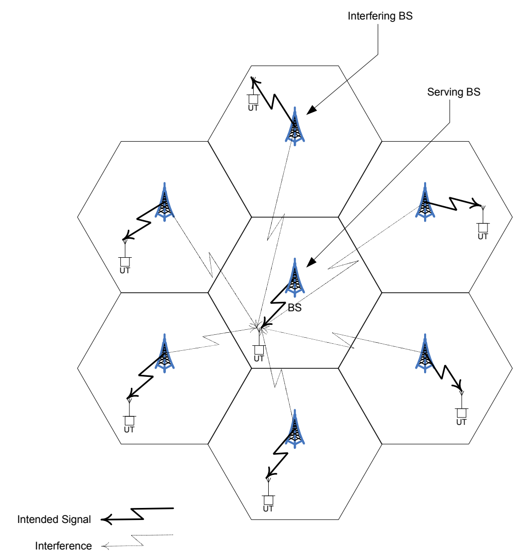

Fig. 1. An interference limited cellular system employing full resource reuse

The signal to interference-plus-noise ratio (SINR) is given by

$$\Gamma_{u_n} = \frac{G_{u_n,n} \mathbb{E}|X_{u_n}|^2}{\mathbb{E}|Z_{u_n}|^2 + \sum_{i \neq n}^{N} G_{u_n,i} \mathbb{E}|X_{u_i}|^2}$$

Denoting the transmit power used by an AP to serve a user u_n by P_{u_n}, we have $\mathbb{E}|X_{u_n}|^2 = P_{u_n}$. We also assume $\mathbb{E}|Z_{u_n}|^2 = \sigma^2$. Note that $G_{u_n,i}$ reflects the composite channel gain, possibly including fast fading.

III. Joint Power Allocation and Scheduling Problem

The joint power allocation and scheduling problem consists of finding the power allocation vector and scheduling vector that will maximize the chosen utility function: network capacity. To facilitate the formulation of the problem, we state the following definitions:

Definition 1: A ***scheduling vector*** $\boldsymbol{U}$ contains the set of users simultaneously scheduled across all cells:

$$\boldsymbol{U} = \begin{bmatrix} u_1 & \cdots & u_n & \cdots & u_N \end{bmatrix},$$

where $[\boldsymbol{U}]_n = u_n$. Noting that $1 \leq u_n \leq U_n$, where U_n is the number of users in cell n, the feasible set of scheduling vectors is given by

$$\mathcal{U} = \{\boldsymbol{U} \mid 1 \leq u_n \leq U_n\}.$$

Definition 2: A ***transmit power vector*** $\boldsymbol{P}$ contains the transmit power values used by each AP to communicate with their respective user:

$$\boldsymbol{P} = \begin{bmatrix} P_{u_1} & \cdots & P_{u_n} & \cdots & P_{u_N} \end{bmatrix},$$

where $[\boldsymbol{P}]_n = P_{u_n}$. Due to a peak power constraint,

$$0 \leq P_{u_n} \leq P_{\max},$$

the feasible set of transmit power vectors is given by

$$\mathcal{P} = \{\boldsymbol{P} \mid 0 \leq P_{u_n} \leq P_{\max}\}.$$

A. System Performance

The SINR for scheduled users will depend on the scheduling vector $\boldsymbol{U}$, as well as on the transmit power vector $\boldsymbol{P}$. We can express the SINR in cell n as

$$\Gamma([\boldsymbol{U}]_n, \boldsymbol{P}) = \frac{G_{u_n,n} P_{u_n}}{\sigma^2 + \sum_{i \neq n}^{N} G_{u_n,i} P_{u_i}}.$$

Assuming an ideal link adaptation protocol and perfect CSI at the transmitter, the per-cell network capacity will be a function of the scheduling vector and transmit power vector. This can be expressed in bits/sec/Hz/cell using the Shannon capacity,

$$\mathcal{C}(\boldsymbol{U}, \boldsymbol{P}) \triangleq \frac{1}{N} \sum_{n=1}^{N} \log\left(1 + \Gamma([\boldsymbol{U}]_n, \boldsymbol{P})\right). \quad (1)$$

B. Optimal Power Allocation and Scheduling

Taking (1) as the objective function we want to maximize, the optimal power allocation and scheduling problem can be formulated as follows:

$$(\boldsymbol{U}^*, \boldsymbol{P}^*) = \arg \max_{\substack{\boldsymbol{U} \in \mathcal{U} \\ \boldsymbol{P} \in \mathcal{P}}} \mathcal{C}(\boldsymbol{U}, \boldsymbol{P}). \quad (2)$$

This consists of finding the optimal scheduling vector $\boldsymbol{U}^*$ and transmit power vector $\boldsymbol{P}^*$ which maximize the network capacity. However, the solution is hard to realize due to the non-convexity of the problem. Indeed, some efforts have been made in this regard by modifying the objective function to make it convex [3], but comprehensive results for this problem still remains an open issue.

One way to simplify this problem is by first adopting a power control policy and then scheduling the users to maximize capacity. Thus, if we assume a received signal-level based power control policy [7], where the AP adjusts its transmit power so that the user attains a target received power, the scheduling vector can be calculated based on this knowledge. If we denote the target received power as R^* then we have the transmit power used to communicate for each user,

$$P_{u_n} = \min\left\{\frac{R^*}{G_{u_n,n} P_{\max}}, 1\right\} \cdot P_{\max},$$

and the new feasible set becomes

$$\mathcal{P}^R = \left\{P_{u_1}, \ldots, P_{u_N} \mid P_{u_n} \in \left[\frac{R^*}{G_{u_n,n}}, P_{\max}\right]\right\}.$$

The problem then transforms to

$$(\boldsymbol{U}^*, \boldsymbol{P}^*) = \arg \max_{\substack{\boldsymbol{U} \in \mathcal{U} \\ \boldsymbol{P} \in \mathcal{P}^R}} \mathcal{C}(\boldsymbol{U}, \boldsymbol{P})$$

For this problem a surprisingly simple and optimal (under the proposed received signal-level policy) solution is obtained [5].

Lemma 1: In a large full reuse network employing a received signal-level power control policy, the optimal scheduling vector consists of users with the highest intra-cell channel gain $G_{u_n,n}$.

Proof: Refer to [5]. ■

We draw the reader's attention however to the fact that extracting the maximum capacity gain in general will involve *joint* optimization of user scheduling and the power control policy. This will be discussed in the next section.

IV. Binary Power Allocation and Scheduling

An interesting result pertaining to problem (2) for the two cell case is presented in [6]. It is shown that in this case the optimal power allocation for any scheduling vector lies in the binary feasible set

$$\mathcal{P}^B = \{P_{u_1}, \ldots, P_{u_N} \mid P_{u_n} = 0 \text{ or } P_{u_n} = P_{\max}\}.$$

Moreover, it is suggested by simulations that with a greater number of cells this *binary allocation* is close to optimal for the majority of the cases. This motivates the adoption of the binary feasible set of transmit powers, $\mathcal{P}^B$ also for an arbitrary number of cells, and allows us to restate the optimization problem as

$$(\boldsymbol{U}^*, \boldsymbol{P}^*) = \arg \max_{\substack{\boldsymbol{U} \in \mathcal{U} \\ \boldsymbol{P} \in \mathcal{P}^B}} \mathcal{C}(\boldsymbol{U}, \boldsymbol{P}) \quad (3)$$

It is interesting to note that this problem can be interpreted as a dynamic version of frequency reuse patterns. A solution to this problem will adjust the reuse pattern by switching on and off cells as required so that system capacity is maximized. We notice that since the optimization variables have a discrete domain, an optimal solution can in principle be obtained through an exhaustive search. However, this proves to be impractical due to the size of the search space: $|\mathcal{U}| = U^N$, where U is the number of users in each of the N cells, and $|\mathcal{P}^B| = (2^N - 1)$.

Intuitively, as the number of users in each cell increases, the probability of finding a user having a high in-cell channel gain, while being sufficiently protected from inter-cell interference, increases. As the number of users increases, this will result in the following scheduling and power allocation vectors, as also later suggested by numerical results in Section V:

$$[\boldsymbol{P}^*]_n = P_{\max}$$

$$[\boldsymbol{U}^*]_n = \arg\max_{u_n} \Gamma(u_n,\, \boldsymbol{P}^*).$$

In this scheme all cells will be on and the user with the best SINR is scheduled; we thus call this scheme MAX-SINR-ON. Note however, that with a small number of users, or an alternate user scheduling policy, capacity gains may be achieved by deactivating cells which do not offer enough capacity to outweigh the degradation caused through interference to the system. In what follows, using an assumption on the interference, we will develope a low-complexity algorithm which determines which (if any) cells are to be deactivated.

A. *Low-Complexity Greedy Algorithm for Power Allocation and Scheduling*

Using results for large idealized networks [5], we can simplify the objective function given in (1) to propose a low-complexity greedy algorithm for choosing which cells (if any) to deactivate.

1) Interference-Ideal Networks: Full spectral reuse has benefits in terms of increased spectral efficiency, but excess interference diminishes the gain associated with increasing reuse. Fortunately, full reuse networks lend themselves to simpler modeling of the total interference experienced by the user, due mostly to the large number of interference sources adding up and averaging at the receiver.

To obtain this simplified model we define the concept of an *interference-ideal* network, as one in which *for any cell user, the total interference received by this user is independent of its location in the cell*. Mathematically, a network with N_{on} active cells is interference-ideal if, for any user u_n and cell $n \in N_{\text{on}}$,

$$\sum_{i \neq n}^{N_{\text{on}}} G_{u_n,i} P_{u_i} = G \sum_{i \neq n}^{N_{\text{on}}} P_{u_i},$$

where G does not depend on the location of u_n. Fortunately, the interference-ideal network is a good model for a full reuse network with a large number of cells. For large N_{on} with all cells transmitting at $P_{\max}$, we have:

$$\sum_{i \neq n}^{N_{\text{on}}} G_{u_n,i} P_{\max} \approx 1/N_{\text{on}} \sum_{i \neq n}^{N_{\text{on}}} G_{u_n,i} \sum_{i \neq n}^{N_{\text{on}}} P_{\max}$$

$$\text{where} \quad 1/N_{\text{on}} \sum_{i \neq n}^{N_{\text{on}}} G_{u_n,i} \approx G, \quad \text{for large } N_{\text{on}},$$

$$\text{thus} \quad \sum_{i \neq n}^{N_{\text{on}}} G_{u_n,i} P_{\max} \approx G(N_{\text{on}} - 1) P_{\max}.$$

Intuitively the approximation above can be explained as follows: No matter where the user is, within his cell, the interference is the sum of all co-channel terms, each of them with a power that depends on the distance from the user to the co-channel APs. Because the user is constrained in his cell and there are a large number of isotropically distributed interferers, the sum power depends *little* on the location of the user within his cell. Moreover, the variation of the interference from the cell center to the boundary in a dense network can be shown to be quite small [5] and from an algorithmic design point of view, G can be considered as the worst case, best case or average out-of-cell channel gain. Note that this model is employed in this paper only to obtain a simplified algorithm, and not for the numerical results: *The numerical results presented in Section V are based on a realistic channel model for all AP-UT links*.

Under this idealized approach, if $\mathcal{N}$ is the set of active cells the objective function from (1) simplifies to

$$\hat{\mathcal{C}} = \frac{1}{N} \sum_{n \in \mathcal{N}} \log \left(1 + \frac{G_{u_n,n} P_{u_n}}{\sigma^2 + (|\mathcal{N}| - 1) G P_{\max}} \right), \quad (4)$$

and the optimization problem becomes,

$$(\hat{\boldsymbol{U}},\, \hat{\boldsymbol{P}}) = \arg \max_{\substack{\boldsymbol{U} \in \mathcal{U} \\ \boldsymbol{P} \in \mathcal{P}^B}} \hat{\mathcal{C}}(\boldsymbol{U},\, \boldsymbol{P}). \quad (5)$$

The procedure for finding an algorithm to solve the problem in (3) is simplified by using (4). Later in Section V numerical results obtained using realistic models verify robustness of these algorithms with respect to our assumption.

From (4) we find that no matter how many APs transmit with full power all users see the same interference from individual APs. Thus, assuming that all APs transmit at $P_{\max}$, complexity can be reduced by employing the maximum SINR scheduler to obtain the scheduling vector $\hat{\boldsymbol{U}}$. Initially, letting

$$[\boldsymbol{P}']_n = P_{\max},$$

then

$$\hat{\boldsymbol{U}} = \arg\max_{u_n} \Gamma(u_n,\, \boldsymbol{P}').$$

As we assume that the interference received by any AP is the same, a search can then be done over all possible transmit power vectors to maximize network capacity. However, even for this algorithm the search over all possible transmit vectors $(2^N - 1)$ still poses a major computational problem. Fortunately, the simplified problem in (5) allows us to avoid searching over all power vectors. We start by keeping all cells on, avoiding unnecessary calculations if there are a sufficient number of users and there is no need to turn off any cells. According to (4), starting with $[\boldsymbol{P}]_n = P_{\max}$ and turning off any cell will result in the same amount of reduction in interference to the other cells. However, the decrease in capacity depends on which cell is turned off. Therefore, the user (cell) with the least SINR (which contributes the least to capacity) can be allocated zero transmit power to see if the gain in network capacity due to reduced interference is more than the capacity lost by turning the cell off. If there is a gain, the cell is allocated 0 in the transmit power vector and the same procedure is repeated with the cell having the least SINR from the remaining cells. When there is a loss in network capacity the procedure is stopped, since turning off any of the remaining cells will always result in a decrease in network capacity. The pseudo-code for the low-complexity greedy approach is detailed in Algorithm 1.

Even though a centralized entity is needed, from an implementation point of view this algorithm significantly reduces the combinations of transmit power vectors that need to be examined from $2^N - 1$ to at worst N. Furthermore, as users are scheduled on the basis of SINR, this avoids having to do calculations over U^N user combinations. Compared to the full exhaustive search required for the problem in (3), this algorithm requires channel knowledge only pertaining to the scheduling vector $\hat{\boldsymbol{U}}$, which translates into significantly less signaling overhead than gathering data on all users. Moreover, there is no need for updating channel information at every iteration (provided that the coherence time is more than the run time of the algorithm).

B. Fairness-Capacity Trade-Off

In resource allocation there is always a trade-off between fairness and capacity. In this work we address a scheme for increasing the system capacity by deactivating cells, so it is obvious that fairness issues will arise. Measures to provide fairness, similar to the case of single-cell scheduling, e.g. by use of proportional-fair type measures [8], can also be considered at the system level. However, going into detail on this issue is beyond the scope of this paper.

Algorithm 1 Low-Complexity Greedy Algorithm

$[\boldsymbol{P}(1)]_n = P_{\max},\ \forall\ n$
for $n = 1 : N$ **do**
  $[\boldsymbol{U}(1)]_n = \arg\max_{u_n} \Gamma(u_n,\, \boldsymbol{P}(1))$
end for
$\mathrm{C}(0) = 0$
for $i = 1 : N$ **do**
  $\mathrm{C}(i) = \mathcal{C}(\boldsymbol{U}(i),\, \boldsymbol{P}(i))$
  if $\mathrm{C}(i) < \mathrm{C}(i-1)$ **then**
    break
  else
    $\hat{\boldsymbol{P}} = \boldsymbol{P}(i)$
    $[\boldsymbol{P}(i+1)]_{\arg\min_n \mathrm{SINR}([\boldsymbol{U}(\mathrm{i})]_\mathrm{n})} = 0$
    remove $[\boldsymbol{U}(i+1)]_{\arg\min_n \mathrm{SINR}([\boldsymbol{U}(\mathrm{i})]_\mathrm{n})}$
  end if
end for

V. Numerical Results

In this section, we present performance results for MAX-SINR-ON and the greedy algorithm based on Monte-Carlo simulations. These algorithms are compared against fixed reuse schemes with cluster size of 3 and 4, in which all cells transmit at $P_{\max}$. In order to ensure a fair comparison, the fixed reuse scheme selects users based on the maximum SINR rule, as do MAX-SINR-ON and the greedy algorithm. A hexagonal cellular system functioning at 1800 MHz is considered, consisting of 19 cells, each with a radius of 500 meters. Gains for *all inter-cell and intra-cell* AP-UT channels are based on the COST-231 [9] path loss model, including log-normal shadowing plus fast fading and antenna gains as well. Log-normal shadowing is a zero mean Gaussian distributed random variable in dB with a standard deviation of σ_X dB. The fast fading is i.i.d. with distribution $\mathcal{CN}(0,\, 1)$. Simulation parameters are detailed in Table I.

We simulate the variation in network capacity as the number of users is increased. Figure 2 shows that MAX-SINR-ON and the greedy algorithm both outperform fixed reuse, due to both having greater reuse. As the number of users increases, all schemes improve due to the multi-user diversity gain. However, MAX-SINR-ON and the greedy algorithm also have a *reuse gain*.

TABLE I

SIMULATION PARAMETERS

Parameter	Value
σ_X	10 dB
Transmit antenna gain	16 dB
Receive antenna gain	6 dB
$P_{\max}$	1 Watt

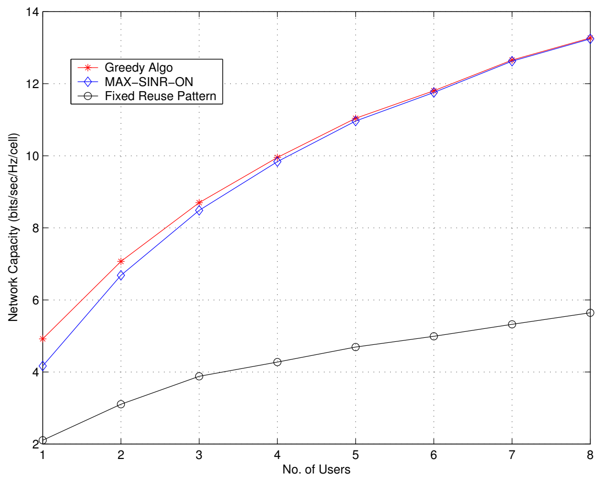

Fig. 2. Network capacity vs. number of users for $N = 19$. Full reuse combined with multi-user diversity provide significant gain as the number of users increases.

When there are only a few users the greedy algorithm outperforms MAX-SINR-ON because it turns off cells to improve the network capacity. Note that for just one user there is no multiuser diversity gain, and therefore this demonstrates the gain associated with the greedy approach for a round robin type of scheduling policy. The network capacity gain of the greedy approach over MAX-SINR-ON is almost 30%. As the number of users increases, both MAX-SINR-ON and the greedy algorithm converge to the same performance due to the maximum SINR scheduling rule, and it is always best to keep all cells on. Figure 3 demonstrates this by showing that even for as few as $U = 8$ users, more than 95% of the cells are kept active by the greedy algorithm.

A conclusive result would be comparing against the optimal scheduling vector and transmit power obtained through an exhaustive search, but Monte-Carlo simulations prove impractical even for a small network (e.g. if $N = 19$ and $U = 8$ then the number of combinations are $(2^N - 1) \cdot U^N = 7.6 \times 10^{22}$).

VI. CONCLUSION

We have addressed the problem of joint power allocation and scheduling for multicell wireless networks.

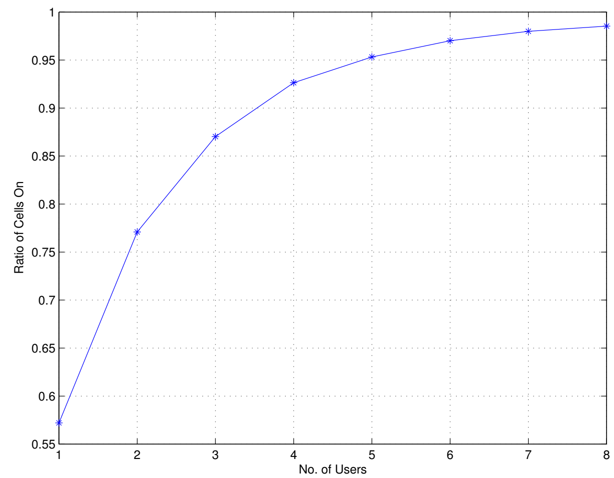

Fig. 3. The ratio of cells which transmit increases as the number of users increases due to the increase in probability of finding a user robust against interference.

We propose a low-complexity greedy algorithm which successively turns off cells as long as the network capacity increases. As the number of users increases, the solution of the greedy approach results in all cells being active and employing maximum SINR scheduling. The proposed scheme exhibits a significantly increased performance compared to having a fixed reuse pattern.

REFERENCES

[1] R. Knopp and P. Humblet, “Information capacity and power control in single-cell multiuser communications,” in *Proc. IEEE ICC*, vol. 1, Seattle, WA, June 1995, pp. 331–335.

[2] M. Andrews, K. Kumaran, K. Ramanan, A. Stoylar, P. Whiting, and R. Vijayakumar, “Providing quality of service over a shared wireless link,” *IEEE Commun. Mag.*, vol. 39, no. 2, pp. 150–154, Feb. 2001.

[3] X. Qiu and K. Chawla, “On the performance of adaptive modulation in cellular systems,” *IEEE Trans. Commun.*, vol. 47, no. 6, pp. 884–895, June 1999.

[4] J. E. Kirkebø, D. Gesbert, and S. G. Kiani, “Maximizing the capacity of wireless networks using multi-cell access schemes,” in *Proc. IEEE SPAWC*, Cannes, France, July 2006.

[5] S. G. Kiani and D. Gesbert, “Optimal and distributed scheduling for multicell capacity maximization,” *IEEE Trans. Wireless Commun.*, Submitted 2006.

[6] A. Gjendemsjø, D. Gesbert, G. E. Øien, and S. G. Kiani, “Optimal power allocation and scheduling for two-cell capacity maximization,” in *Proc. IEEE RAWNET (WiOpt)*, Boston, MA, Apr. 2006.

[7] J. F. Whitehead, “Signal-level-based dynamic power control for cochannel interference management,” in *Proc. IEEE VTC*, Secaucus, NJ, May 1993, pp. 499–502.

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

[9] E. Damosso and L. M. Correia, Eds., *Digital Mobile Radio Towards Future Generation Systems*. Brussels, Belgium: European Commision, 1999, COST 231 Final Report.