

Dynamic Load Balancing Through Coordinated Scheduling in Packet Data Systems

Suman Das, Harish Viswanathan, Gee Rittenhouse
Wireless Technology Research
Lucent Technologies, Bell Labs.
600 Mountain Ave, Murray Hill, NJ 07974.
{sumand, harishv, gee}@lucent.com

Abstract—Third generation code-division multiple access (CDMA) systems propose to provide packet data service through a high speed shared channel with intelligent and fast scheduling at the base-stations. In the current approach base-stations schedule independently of other base-stations. We consider scheduling schemes in which scheduling decisions are made jointly for a cluster of cells thereby enhancing performance through interference avoidance and dynamic load balancing. We consider algorithms that assume complete knowledge of the channel quality information from each of the base-stations to the terminals at the centralized scheduler as well as a two-tier scheduling strategy that assumes only the knowledge of the long term channel conditions at the centralized scheduler. We demonstrate that in the case of asymmetric traffic distribution, where load imbalance is most pronounced, significant throughput gains can be obtained while the gains in the symmetric case are modest. Since the load balancing is achieved through centralized scheduling, our scheme can adapt to time-varying traffic patterns dynamically.

I. INTRODUCTION

Third generation cellular systems are currently being designed with the emphasis of providing high speed packet data service on the downlink through a shared high data rate channel. The base-station arbitrates this channel among the different users based on channel condition feedback and backlog [1], [2], [3]. Variable rate transmission at fixed power is a key characteristic of these systems. In the current approach, the scheduler resides at the base-station and operates independently across different cells. In other words, a cluster of contiguous base-stations transmit signals to different users independent of the activity or the scheduling decisions made at the neighboring base-stations. We consider a system in which a cluster of neighboring base-stations schedule transmissions to users in their cells jointly. We refer to this as *coordinated scheduling*. The coordinated scheduling can be implemented through a centralized scheduler entity or through signaling between the base-stations on the backhaul network. With coordinated scheduling we can potentially incorporate interference avoidance to some degree (depending on the cluster size) by turning off some base-stations in the cluster while other base-stations are transmitting to users at the edges of their cells. Furthermore, the coordinated scheduling can adapt to traffic imbalances dynamically since any base-station within the cluster can transmit to any user within the cluster. We propose some coordinated scheduling algorithms

and evaluate their potential gains in this paper. Our results show that for asymmetric traffic scenarios significant gains are obtained through load balancing from coordinated scheduling. We do not consider phase-coherent transmission of radio-frequency signals across base-stations or joint encoding of the information signals across base-stations. These would require considerably more coordination between base-stations but can potentially provide more gains.

Load balancing has been considered in the literature for both time division and code division multiple access networks. Traditional time division multiple access (TDMA) cellular networks employed frequency reuse to guarantee sufficient signal-to-interference-and-noise-ratio (SINR) at the receiver for transmission of information. At the time of deployment the available spectrum is assigned to the different cells according to the desired frequency reuse. A major disadvantage of such networks with fixed frequency assignment was the inability to adapt to asymmetries in traffic patterns that could potentially be dynamic. As a consequence dynamic channel allocation (DCA) has been studied extensively in the literature for such TDMA voice networks [4], [5]. With the advent of code-division multiple access (CDMA) technology frequency planning was no longer required because of universal reuse. Nevertheless, one could consider designing the network to match the traffic conditions at the time of deployment through base-station location, antenna down-tilt and base station transmit power [6], [7], [8]. To date most of these optimization techniques have been applied to optimize the network for non-shared channels. Furthermore, these optimization procedures are performed only at the time of deployment and hence are not dynamic. The coordinated scheduling considered in this paper provides a method for dynamic load balancing in packet data CDMA networks. Inter-cell coordination with variable power and variable rate transmission was considered in [9]. However, they did not consider scheduling in combination with inter-cell coordination.

The main contributions in this paper are the proposed algorithms for coordinated scheduling, characterizing their potential gains and understanding of the impact of channel fading on the gains. The paper is organized as follows. A detailed description of the different system architectures considered in this paper that makes it possible to do coordinated scheduling are described in Section II. The coordinated scheduling algo-

gorithms are presented in Section III and the simulations results are presented in Section IV. We then present an algorithm with limited coordination for achieving only the load balancing part of gain from coordinated scheduling in Section V. We conclude in Section VI.

II. SYSTEM ARCHITECTURE

We consider a multi-cell wireless cellular network consisting of numerous base-stations. Mobile terminals are randomly distributed in the service area and at any given instant, each base-station serves multiple users. We consider a code division multiple access (CDMA) like system with universal frequency reuse where the entire spectrum is used in every cell and the uplink and downlink are assigned separate spectrum. We restrict our attention on the best effort packet data service on the downlink and assume for simplicity that the entire spectrum and base-station power is devoted to this service as in the case of 1X-DO [10] system. We focus our attention on the scheduling or the resource allocation problem for the downlink channel. In a typical third generation CDMA packet data system based on a single high rate shared channel, time is divided into multiple time slots and during each time slot a single user is selected by the scheduler for transmission based on the backlog and the channel conditions fed back by the users on the uplink. During each time slot the entire transmit power from a base and all the CDMA codes are assigned to a single user. Several scheduling algorithms have been proposed in the literature that meet different objectives [2], [11], [12], [14]. All of these scheduling algorithms are designed under the assumption that each base station schedules transmissions independently following the system architecture described below.

Before we describe the scheduler architectures we define the notion of *load* for packet data service. In the absence of any quality of service guarantees to the users, the concept of load is not well defined since the service to each user can become arbitrarily poor with increasing number of users. Thus some notion of quality of service is necessary to specify the load in the system. Since we are concerned with best effort packet data service we do not impose any delay guarantees on the service provided. Instead we use the long term average throughput achieved by the users as the quality of service metric and require that all users are treated the same. To this end we consider token queue based scheduling algorithms in which for each user a token queue is maintained [14]. The token queues are incremented at every slot according to a given arrival rate that is same for all users and decremented according to their respective services rates. With this approach the load in the system is completely determined once the number of users and their location within the cell is specified. The scheduling algorithms are described in more detail in Section III. Note that the equal token rate for all users is considered only for simplicity and can be easily modified to different token rates for different users to provide unequal service.

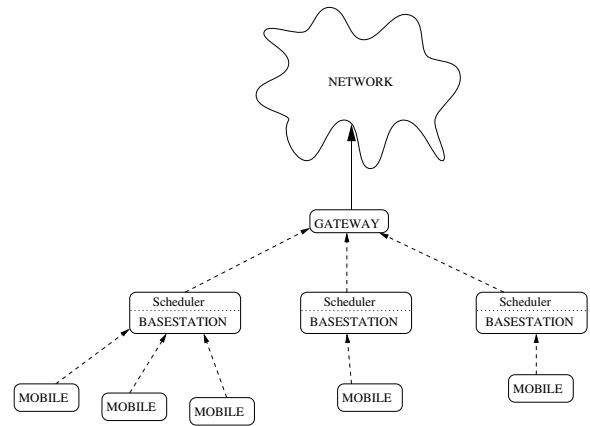


Fig. 1. Architecture for baseline scheduler

A. Independent Scheduler Architecture

The conventional architecture, that serves as our baseline for performance comparisons, is illustrated in Figure 1. The gateway sits at the edge of wireless service provider network and the Internet and is responsible for encapsulating the network layer packets into link layer packets according to the link protocol, such as the point-to-point protocol (PPP), used in the service provider network so as to easily provide mobility. The gateway directs the packets intended for a user to the appropriate base-station which maintains the wireless link to the user. While in practice a real network could have multiple intermediate nodes between the gateway and the base-station, for our purposes it is sufficient to abstract the network as shown in the figure. The primary physical layer and MAC layer scheduling functionalities are handled by the base-station operating independently of other base-stations.

In this architecture (Figure 1) each mobile is assigned to a particular primary base-station that is responsible for transmission of data from the network to the user. The base-station scheduler, based on the channel state information and the backlog, identifies the user to be served in every time-slot independently of which users in the neighboring cells were selected to be served by the neighboring base-stations.

Typically in such an architecture the association between the mobile-station and the corresponding base-station is mobile directed. Each mobile listens to the pilot signal transmitted from all the surrounding base-stations and measures the channel quality, quantified through pilot carrier to interference (C/I) ratio, from each of these base-stations and identifies the single best base-station that has the best channel quality. It then associates itself to this best base-station to receive data from and transmits the channel quality of this link on the uplink. The channel quality and the available power at the base-station determine the maximum data rate that can be assigned to the particular user. Thus, in the baseline architecture the base-station selects between the available users that are associated to it based on their current channel conditions and possibly their data queue size. There is little cooperation among the base-stations to control the interference or share the load.

It should be noted that the channel condition between users and the base-stations is not static. Channel conditions depend on the distance between the base-stations and the mobiles, long term fading like the shadow fading and short term fading due to relative motions between the users and the base-stations [4]. The short term fading is typically on the time scale of a few slots. In the baseline architecture users can potentially associate themselves to two different base-stations in two consecutive time slots assuming ideal rapid signaling. However, this is unlikely in practice. Even though base-stations can have instantaneous channel knowledge, handing over users from one base-station to another involves ripping down connection between the gateway or some other intermediate node (one hop away from the base-stations) and the existing base-station and reestablishing a channel with the new base-station. Moreover, unsent data needs to be forwarded to the new base-station. For the purposes of comparison with our proposed algorithms, we consider the extreme cases within this architecture, one in which the mobiles associate to the same base-station over a long time duration and the other case in which users can change their association to the base on every time slot. In the *static scheduling* scheme mobiles associate to the base-station based on the long term channel conditions that include the path loss and the shadow fading.

The second scheduling policy, called the *Fast Cell Site Selection (FCSS)* allows a mobile to select the best base-station based on instantaneous channel conditions, including the fast fading and hence the association can change from one time slot to the next. Note that in this architecture with mobiles deciding which base-station to receive data from, it is not possible to load balance dynamically since the mobiles are unaware of the load at the base-stations. In the next subsection we propose an architecture in which we move the scheduler to a higher level in the hierarchy to facilitate load balancing.

B. Centralized Scheduler Architecture




Fig. 2. Architecture for centralized scheduler

Actual wireless systems contain many base-stations and hence the signal received by a mobile from the serving base-

station is corrupted by the signals transmitted by all other base-stations. In general therefore it is suboptimal for each base-station to allocate user resources independently of other base-stations. Ideally we should consider activities of all the base-stations and control them simultaneously in order to maximize the overall system performance. This means that we should consider the channel conditions from all the base-stations to all the users and also all their backlog (queue size) states and optimize the resource allocation. It is obvious that for a large scale wireless network it is impractical to include all the base-stations in the optimization. The payload over the backhaul network that gathers the channel information to the centralized scheduler would be enormous. Thus for a realistic solution we divide the entire wireless network into several clusters (Figure 2). The clusters need not be physically disjoint and different sectors of a single base-station can belong to different clusters as explained in later sections. Each such cluster comprises of multiple base-stations. The scheduling functionality is now performed for the entire cluster. Thus the scheduling functionality is moved to one hop away from the base-station and no longer resides at the base-station. In order to facilitate centralized scheduling within the cluster we require that the mobiles send back the channel quality information on the links from all the base-stations in the cluster on the uplink channel. This can be done for example by time multiplexing the channel quality of the different links thereby using the same number of uplink bits but providing fewer updates. The base-stations then forward this information to the centralized scheduler. The base-stations also forward to the centralized scheduler the status of the buffers for the different users. Alternately, the queues can be maintained at the scheduler node itself and forwarded to the appropriate base-station for transmission to the users.

Scheduling now involves determining which base-stations in the cluster should transmit in each time slot and to which users they should transmit data to. The association of the users to the base stations is now determined by the centralized scheduler and hence makes it possible to achieve load balancing since the scheduler is now aware of the queue sizes at each of the base-stations in the cluster. With the centralized scheduler it is also possible to prevent neighboring base-stations from transmitting when a particular user near the edge of a cell is being served. Thus we expect that this mechanism of interference avoidance would provide gains in addition to load balancing that is naturally achieved.

When the instantaneous channel conditions from all base-stations in cluster to each of the mobiles are available to the centralized scheduler we call the corresponding scheduling policy the *coordinated scheduling* policy. However, this policy demands a high bandwidth connection between the base-stations in the cluster and the scheduler. Alternatively, we can consider a *two-tier scheduler*. In this approach, only the long term average channel conditions are assumed available to the central scheduling entity. This information is used to determine which of the base-stations will transmit in the next time-slot, while the instantaneous data rate from the transmitting base-

stations is determined by the instantaneous channel values available only at the base-stations. We continue to assume that the mobiles feed back instantaneous channel conditions from all the base-stations. However, only the long term channel condition and backlog status is sent to the central scheduler and hence reduces the backhaul bandwidth requirements. Note that the users can still be served by different base-stations in subsequent time slots as in FCSS and hence maintaining the buffers at the centralized scheduler node is more favorable. The detailed description of the algorithms are presented in the next section. In Section V, we present a dynamic load balancing algorithm within this architecture that further reduces the requirements by allowing only slow migration of users across the base-stations.

III. SCHEDULING ALGORITHMS

We model the received signal for user k in the network by

$$r_k(t) = h_i(t)\sqrt{P_{k,i}}s_i(t) + \sum_{j=1, j \neq i}^N h_j(t)\sqrt{P_{k,j}}s_j(t) + n(t) \quad (1)$$

where $s_i(t)$ is desired information signal received from the base station i transmitting to this mobile and is received at averaged received power level $P_{k,i}$, $s_j(t)$ is interference signal received from the interfering base-station $j \neq i$ at average power level $P_{k,j}$, $h_j(t)$ is the fast fading gain from any base j and $n(t)$ is the additive Gaussian noise at the receiver with spectral density N_0 . The shadow fading is incorporated into the average received powers P_k . The slow shadow fading is modeled as independent lognormal variables while the fast fading is modeled by Rayleigh and Rician distributions. We model a total of N base-stations in the network in a hexagonal layout that are divided into multiple clusters with N_c base-stations in each cluster. We assume that each of the base-stations, when they transmit, transmit with the same power P_{max} to a single user in any given time-slot. The signal from the base-station to the user decays with the distance according to a simple power law with a decay exponent of 3.5.

The signal received by any user from all the base-stations except the one that is serving the user is treated as interference. For all the users in the cluster we calculate the interference signal from all the interfering base-stations. For interference calculation we consider the central cluster is surrounded by two rings of base-stations. Thus the interference power from each interfering base-station j is $P_k(j) = \frac{cP_{max}}{d_j^{3.5}}$, where d_j is the distance between the base-station j and the user. The constant c corresponds to the intercept in the path loss model and is assumed to be 28.5 dB when the distance is in meters [15].

Thus the total interference plus noise power to users under the base-station 1, is given by

$$I = \sum_{j=2}^N \left(\frac{cP}{d_j^{3.5}} \right) + N_0. \quad (2)$$

It should be mentioned here that since this is a CDMA system all interference terms are from co-channel interference.

Moreover we assume all the codes are used in this system and hence the processing gain is assumed to be 1. For convenience, we make a distinction between two types of interfering signals. The interference from all the base-stations that are not in the cluster are calculated as above while for interfering base-stations within the cluster we also take into account the fading components. This is justified since the number of base-stations in the cluster N_c is typically small while the total number of interfering base-stations is large and hence the effect of the fast fading of the signals from all the base-stations on the total interference power can simply be approximated by the average received power. When a user is selected for transmission then the data rate to the user is assumed to be given by the Shannon limit $R = \log_2(1 + SINR)$ where $SINR$ represents the received signal to interference ratio from the base at the mobile.

We assume that all users in the cluster under consideration are backlogged but use token queues to guarantee fairness among the users. A token queue is simply a counter that is incremented at a rate slightly above the target throughput rate in every slot and decremented according to the service rate. In the scheduling policy considered, the base-stations try to provide uniform and equal long term data rate to all the users by setting the token rate equal for all users. Such a scheduling policy is considered as an example scenario. The coordinated scheduling algorithm proposed can be generalized to other scheduling policies with different objectives since most of these algorithms fit into the framework maximizing a weighted sum of rates.

In the following subsections we provide details about four different algorithms that we consider. The first two are based on the independent scheduler architecture while the last two are based on the centralized and the two-tier architectures.

A. Static Scheduler

This algorithm is based on the Independent Scheduler Architecture described in the previous section and is a static scheduler in the sense that it employs an independent scheduler at each base-station and assigns users to base-station statically based on the long term average signal. This algorithm needs the least amount of information to be fed back to the scheduler. We first instantiate the users by placing them randomly within the cells.

At the beginning of each assignment stage

- Measure the interference (I_N) from all base-stations that are outside the cluster using (2) assuming that they transmit at full power.
- For all the base-stations within the cluster, the average signal strength (that includes the effect of the shadow fade modeled as lognormal in the simulations) is measured per base. Denote these signals by $\overline{S}_{(b,k)}$, $\forall b \in \{1 \dots N_c\}$ and users $1 \leq k \leq K$.
- Assign user k to base-station b^* from which it receives the strongest average signal $b^* = \arg \max_{b \in \{1, \dots, N_c\}} \overline{S}_{(b,k)}$

Let users $\{1, \dots, K_1\}$ be assigned to base-station 1 in the cluster. Note that users are assigned to only one base-station

and the assignment remains the same over a large number of time slots since it does not depend on the fast fading. For each user a token queue is maintained and updated after each time slot to guarantee fairness. Let the token queue depth of these users be $\{q_1, q_2, \dots\}$ and let the token arrival rate be λ .

- The instantaneous received signal strength $S_{i,k}(t)$ is given by the product of the average signal strength $S_{(1,k)}$ and the fast fading component: $S_{i,k}(t) = \overline{S_{(1,k)}} |h_{i,k}(t)|^2$.
- The instantaneous SINR is given by

$$SINR_{(1,k)}(t) = \frac{S_{(1,k)}(t)}{\alpha S_{(1,k)}(t) + \sum_{i=2}^{N_c} S_{i,k}(t) + I_N + N_0},$$

where α is the self interference factor that models the transmitter and receiver non-linearities and limits the maximum SINR. For our simulation we assume α to be 0.01 which corresponds to the maximum SINR value of 20 dB. The instantaneous allowed rate is calculated as $R_{(1,k)}(t) = \log_2(1 + SINR_{(1,k)}(t))$.

- For the time slot t the scheduler chooses the user k^* that satisfies $k^* = \arg \max_{k \in \{1, \dots, K\}} w_k(t) R_{(1,k)}(t)$, where $w_k(t) = \max(0, q_k(t))$. It is to be noted that we allow queue lengths to assume negative numbers since these are token queues.
- The queues are updated as follows: $q_{k^*}(t+1) = q_{k^*}(t) + \lambda - R_{(1,k^*)}(t)$ and $q_k(t+1) = q_k(t) + \lambda$ for the rest.

B. FCSS

This scheme is also based on the Independent Scheduler Architecture but uses a dynamic cell selection algorithm. Unlike the previous scheme, users are assigned to the base-station from which it receives the strongest instantaneous signal. As before the instantaneous signal strengths are calculated, but now at *each time slot* the users are assigned to the base-station with the largest signal strength. It is worth reiterating that the average signal strength varies only when the user changes its location, but the fast fading channel component still contributes to the instantaneous received signal. Thus the latter can differ from the average signal strength. After users are associated to the base-stations based on their instantaneous signal strength, similar scheduling algorithm as before is used at each base-station. Due to change in instantaneous channel conditions users can move frequently across the base-stations, especially those that lie on the boundary between two cells.

C. Coordinated Scheduler

This algorithm is based on the centralized scheduling architecture and assumes knowledge of instantaneous channel conditions of all the base-station to user links within the cluster at the scheduler. In this algorithm $S_{(i,k)}(t)$, I_N and noise values are calculated as before for each time slot. However, an additional decision is made regarding which base-stations should be kept active at each time slot by the scheduler. The purpose is to control the interference that the users experience from base-stations within the cluster. There are two contradictory forces at work. The activities of neighboring base-stations introduce extra interference and thereby reduce

the rate at which data can be transmitted to a particular user. This is especially true for regions of the cluster that is interference limited, namely the edges of the cells. On the other hand, since we mandate that at any instant of time a base-station can transmit to at most one user, the higher the number of base-stations that are active, higher the number of users that are served. Thus the optimization problem involves finding the right set of base-stations to be active and the appropriate users they should transmit to. Let $\mathcal{A} \subset \{1, \dots, N_c\}$ be the optimum set of active base-stations, then the optimization problem is to find the users that maximize

$$\sum_{b \in \mathcal{A}, k \in K} w_k(t) R_{k,b}^A(t), \quad (3)$$

where $R_{k,b}^A(t)$ is given by

$$\log_2 \left(1 + \frac{S_{(b,k)}(t)}{\alpha S_{(b,k)}(t) + \sum_{a \in \mathcal{A}, a \neq b} S_{(a,k)}(t) + I_N + N_0} \right).$$

The algorithm is a generalization of the maximum-weight scheduling algorithm that is fairly general in terms of achieving different objectives. In particular, if the weights are set according to the token queues then the scheduling algorithm is stability optimal subject to the rules of inter-cell coordination assumed here [13].

Which user should be connected to which base-station actually depends on the set of base-stations that are active. However, even though at first sight it might appear that we may have to consider all the $2^{K N_c}$ possible base-station user combinations, that is not the case. Since we mandate that all base-stations transmit at either full power or remain silent completely, the interference seen by user k from base b is independent of the user that base b is actually transmitting to. The coordinated scheduler takes into account the above insight to reduce the complexity of the algorithm to order of 2^{N_c} combinations as explained below.

- Compute I_N and $S_{(b,k)}(t)$ as before.
- For all $2^{N_c} - 1$ possible active base-station combination (note at least one base-station should be active), we compute the corresponding $SINR_{(b,k)}(t)$ values where b is in the active base-station list.
- For each such configuration, for each active base-station we identify the user with best $w_k(t) \log_2(1 + SINR_{(b,k)}(t))$ value.
- The configuration with the highest objective function given by (3) is determined and for each of the base-user pair the allotted rate $R_{(b,k)}(t)$ is computed.
- The queues are updated as before.

Note that users are not associated to any base-station a priori and the scheduler determines which base-stations transmit and to which users, it is possible to achieve load balancing from asymmetric loading of the cells. This effect is demonstrated in our simulation results.

D. Two-tier scheduler

The coordinated scheduler described above requires knowledge of all possible instantaneous base-station to user channel

conditions to be gathered at a central place. As mentioned in the previous section it is possible that this requirement might put too much burden on the backhaul. Furthermore, the associated delays could be significant making it impractical. Hence we propose the two-tier scheduler described below.

The two stage scheduler, identifies which set of base-stations should remain active based on long-term average signal strength but with updated queue states. Thus the central scheduler only needs periodic channel state update and continuous snapshot of user queue states to determine \mathcal{A} that maximizes $\sum_{b \in \mathcal{A}, k \in K} w_k(t) R_{k,b}^A(t)$ where $R_{k,b}^A(t)$ is given by

$$\log_2 \left(1 + \frac{\overline{S_{(b,k)}}}{\alpha \overline{S_{(b,k)}} + \sum_{a \in \mathcal{A}, a \neq b} \overline{S_{(a,k)}} + I_N + N_0} \right).$$

Once the optimum set of base-stations are determined the instantaneous channel data available to the individual base-stations are used to identify the users that maximize the objective function. In this way the centralized scheduler determines the active base-station set and each individual base-station scheduler selects the particular user. Note that we continue to assume that each base-station itself is aware of the channel conditions to all the users in the cluster. Thus the users can be served by any base-station in each time slot depending on the load at the different base-stations implying that load balancing can be achieved through this scheme as well.

IV. SIMULATION RESULTS

We assume hexagonal cells of radius 1 km. We assume independent 4.3 dB lognormal shadow fading, appropriate for small cells [4]. We consider 3 different types of fast fading, pure Rayleigh fading and Rician fading models with $K = 0dB$ and $K = 6dB$. For each instantiation of user locations and shadow fades, we simulate 2000 time slots. We assume that users do not move within this 2000 time slots. 40 such instantiations for each traffic and fading model combination is generated to obtain the cumulative distribution functions of the average throughputs.

A. Results

We consider a seven cell cluster with each hexagonal cell divided into three non overlapping 120 degree sectors. All the three sectors of the central cell and one sector each from the surrounding cells form a cluster. The users that are geographically located in the cluster are considered for the coordination scheme. We consider several traffic distributions, namely, uniform traffic, traffic where 50% and 99% of the traffic are in one cell.

The performance of the four different algorithms are presented in Figures 3 and 4 for different traffic densities and fast fading scenarios. In each figure there are 6 subplots. The left hand side plots show the results for five users per sector and the right hand side plots are for seventeen users per sector. There are three sets of subplots for each traffic density corresponding to Rayleigh fading and 0dB and 6dB Rician fading schemes, respectively. For each of these cases we plot the cumulative distribution function of the achieved throughput for each of

the algorithms. Since we deploy a token queue scheduling policy to ensure long term uniform throughput the maximum achievable throughput for seventeen user system is lower than that of five user system. But the token queue system cannot guarantee the data rate for certain users that have really bad channel conditions. The effectiveness of a scheduling policy is determined by what percentage of users achieve close to target throughput.

In general we see there is substantial improvement in achievable throughput in using coordinated scheduling over the baseline scheduler for all traffic and channel (fast fading) situations. The improvements are moderate for uniform and dense traffic with Rayleigh fading channels (Figure 3), whereas the gains are substantial for highly asymmetric traffic (Figure 4). Moreover the relative performance of FCSS algorithm over the coordinated scheduling scheme varies with the traffic density and channel conditions.

Even though the performance for a single user in a fading channel is significantly worse than that in Gaussian channel, for a multiuser system one can exploit channel fading to improve the throughput of the system [12]. This gain increases with the increase in the number of users and intensity of channel fades. FCSS improves over the baseline algorithm by exploiting this diversity gain. While users are statically assigned for a long-term in the baseline system, in FCSS they can move across the cells more frequently. So a user that is statically assigned to a particular base-station and going through deep short term fades, may have a better channel from a different base-station and exploit the superior channel condition with FCSS policy. However, as the number of users in a sector increases there are more users to choose from for a particular base-station. So even though there might be some users, statically assigned to a base-station, that are going through deep fades, there is a higher probability that there is another user also statically assigned to the same base-station but seeing better channel conditions. This is reflected in the relative performance of systems with larger (17) number of users over systems that are less densely populated. As the fading variance decreases, from a purely Rayleigh fading channel to a Rician channel with increasing proportion of line-of-sight component, the diversity gain starts diminishing.

The performance of FCSS systems is similar to that of a coordinated scheduler especially in systems with uniform and dense traffic and significant channel fading. The coordinated scheduler also exploits the multiuser diversity as in the case of FCSS policy. Furthermore, it tries to coordinate the scheduling across base-stations. Sometimes the coordinated scheduler might decide to turn off a certain set of base-stations in the cluster especially to serve users that are at cell edges. However, for large number of users distributed uniformly across the entire cluster and for Rayleigh channels such occurrences are rare. The interference reduction gain for edge of the cell users is offset by losing out the opportunity to transmit to another user. In fact for uniform traffic, and especially for Rayleigh channels it is almost always advisable to have all the base-stations transmit simultaneously. The situation changes a bit

for Rician channels especially if the number of users in the system is low.

For asymmetric traffic conditions (Figures 4) however, there is gain in coordinating across base-stations. In these scenarios we see gains beyond simple diversity gain. The coordinated scheduler can take advantage of its inherent load balancing capability. The FCSS scheme blindly associates user to base-stations irrespective of the load of the base-station. The coordinating scheduler can more intelligently assign users to base-stations that have less load. Thus, while we see the FCSS and the baseline scheme are unable to maintain the required throughput guarantee for all users in asymmetric traffic conditions, the coordinated scheduler can do a much better job.

The performance of the simple two-tier scheduler is universally better than the baseline scheduler, but it is not always better than the FCSS system. This is because the decision to turn a base-station on or off in case of 2-tier schedulers are made at the centralized scheduler in the two-tier scheduler based on the long-term average channel conditions and the current snapshot of queues. In case of uniform traffic distribution and Rayleigh fading channels, as mentioned before it is probably almost always better to have all the base-stations on. The two-tier scheduler might decide to turn off certain base-stations based on the backlog information but, this may not be the right decision. In such conditions, (uniform traffic density and Rayleigh fading channels) the instantaneous channel condition is more important than the backlog states. For Rician channels the long-term channel more closely represents the current channel conditions and in those cases these coordination decisions are more accurate. There is still a substantial performance gap between the coordinated scheduler and the two tier scheduler. In Section V we propose another 2-tier scheduler that requires limited coordination with performance closer to the coordinated scheduler under mild fading conditions.

B. General Conclusions

To summarize, there are three types of gain that we can achieve over baseline system:

- Selection diversity gain by appropriately choosing the base-station for a particular user based on instantaneous channel conditions and multi-user diversity gain by base-stations scheduling the users with good channel fades to serve.
- Interference avoidance gain by adjusting interference levels by turning on or off certain base-stations. These should particularly help out the edge users (interference limited users)
- Load balancing by base-stations serving users from neighboring cells whenever possible.

FCSS addresses the first type of gain while our new algorithms try to achieve all the three gains. The simulation results show that,

- 1) For uniform traffic with Rayleigh fading, diversity gains dominate. Thus FCSS scheduling is almost as good

as the more sophisticated coordinated algorithms, especially when there are sufficiently large number of users. For small number of users we still see some gains because of enhanced multi-user diversity gain.

- 2) Under more moderate fading conditions as in Rician fading channel we start seeing the effects of coordinated scheduling.
- 3) Load balancing shows significant gain for asymmetric traffic densities.
- 4) The two-tier scheduler for uniform and Rayleigh channels is actually inferior to FCSS systems. This points to the fact that in those cases instantaneous channel knowledge is very important. The extra interference seen by each user by keeping all base-stations active at all time is offset by the diversity gains.

V. LOAD BALANCING WITH LIMITED COORDINATION

The aforementioned scheduling algorithms for load balancing allow the migration of users across base-stations at every time slots. This could be difficult to achieve in practice because of administrative overhead. While it is necessary to select the best base-station at the time scale of fade durations for obtaining full diversity gain, it is not necessary for the purpose of load balancing because the traffic distribution typically changes slowly. In this section we propose a new scheduling algorithm that requires only limited coordination between the base-stations in the cluster and where users switch base-stations only once for a large number of slots. This algorithm does not require significant amount of data to be shared across the cluster of base-stations and yet performs similar to a centralized scheduler for the purposes of load balancing. The architecture is similar to that of a two-tier scheduler where we assume that the central scheduler will have knowledge of average channel characteristics of all the users to all the bases but not the instantaneous channel knowledge. Unlike the two-tier scheduler in Section III-D in this scheme we allow all the base-stations to be active all the time and hence no scheduling decisions are made at the the centralized scheduler on a time slot basis. Thus the primary aim of the scheduling policy is to distribute the users among the base-stations that achieve load balancing and will not achieve gains from interference avoidance or diversity. Let λ_i be the data arrival rate of user i (or the token rate in the case of minimum throughput guarantees) and $\overline{R}_{b,i}$ the average transmission rate from base-station b to user i under the average channel condition. If U_b be the set of users that are served by the base-station b , then for a static channel for stability, $\sum_{i \in U_b} \frac{\lambda_i}{\overline{R}_{b,i}} < 1$.

However for a fading channel, multiuser diversity can be exploited and more users can be supported. For a stable system we want for all base-stations in the cluster to satisfy the above relationship. In general for proper load balancing we want for any pair of base-stations b and b' ,

$$\sum_{i \in U_b} \frac{\lambda_i}{\overline{R}_{b,i}} \approx \sum_{j \in U_{b'}} \frac{\lambda_j}{\overline{R}_{b',j}}$$

The following steps are performed at the beginning of each assignment period by the centralized scheduler:

- Measure the average channel conditions and hence the average channel data rate $\overline{R}_{b,i}$ of all users to all bases.
- Assign user i to base b^* such that $b^* = \arg \max_b \overline{R}_{b,i}$.
- Let U_b be the set of users assigned to base b and $Q_b = \sum_{i \in U_b} \frac{\lambda_i}{\overline{R}_{b,i}}$.
- Let μ and σ be the mean and standard deviation of Q_b . We assign threshold $T = \mu + \beta\sigma$, where β is an arbitrary constant.
- Create a complete sorted list of the average rates of all the users to all the base-stations.
- For all base-stations initialize $q_b = 0$. Continue the following iterations till the list is exhausted.
 - Select the next best user rate $\overline{R}_{b,i}$ from the completely sorted list.
 - If the user i is not yet assigned to any base, and $q_b + \frac{\lambda_i}{\overline{R}_{b,i}} < T$, then assign user i to base b and set $q_b = q_b + \frac{\lambda_i}{\overline{R}_{b,i}}$.

For the users that are not assigned to any base-station by the above scheme, assign them to the base-station that has the best $\lambda_i/\overline{R}_{b,i}$ value. One iteration of this algorithm may not put all the base-stations' q_b -s under T and we should refine T and loop again, but in our solution we stop after one iteration. Once the users are assigned to base-stations they are served similar to that of a baseline system in which each base-station schedules transmissions to its users independently of other base-stations. Thus the centralized scheduler functionality is required only once every few seconds when the assignment is made.

A. Pilot Power Based Load Balancing

Our proposed coordination schemes mandate that scheduling and handoff decisions are taken by the base-stations. Under current CDMA guidelines, handoffs are mobile directed. One way to circumvent this standard restriction is to reduce the pilot power of the heavily loaded base-stations, without reducing the signal power. Since the decision of associating to a base is taken by the mobile based on the pilot power, one can effectively achieve load-balancing by controlling the pilot power.

B. Simulation Results

We simulate a dynamic scenario where the users continuously change locations. We simulate scenarios where users move at three different speeds, 3 Km/h, 45 Km/h and 99 Km/h. The channel is modeled as a Rician channel with K factor 0 dB and the Rayleigh component modeled by a Jakes model [4]. We assume a carrier frequency to be 2GHz. For the baseline and the load balancing schedulers, users are reassigned to the base-station after every 10 seconds.

We ignore any long term shadow fading effects in this simulation since we wish to demonstrate the load balancing capability. The results are shown in Fig 5. There are 3 different traffic densities. The top row correspond to uniform traffic

density, the middle row represents situation where 50% of the total traffic is in one cell and the bottom set of figures represent scenario where 99% of the users are in the central cell.

Our results show both the load-balancing and the coordinated scheduling algorithms adjust themselves well with dynamic traffic conditions. In skewed traffic scenario both the load-balancing algorithms improve upon the baseline scheduling algorithms.

VI. SUMMARY

Coordinated scheduling across a cluster of base-stations was proposed to improve the performance of high speed packet data service using a shared channel. Coordinated scheduling could potentially provide diversity, interference avoidance and load balancing gains. In symmetric traffic scenarios under Rayleigh fading the diversity gain is the dominant factor and thus fast cell-site selection is sufficient to realize most of the gains. Under less severe fading conditions such as Rician fading with $K > 0$ dB, the other gains become important. With asymmetric traffic and Rician fading we see the maximum gains from coordinated scheduling. We also proposed an algorithm with limited coordination which is capable of realizing load balancing gains under asymmetric traffic conditions. This algorithm requires only periodic coordination using the back haul network and is relatively easier to implement.

REFERENCES

- [1] P. Bender et. al., "CDMA/HDR: A bandwidth-efficient high speed wireless data service for nomadic users", *IEEE Communications Magazine*, 38(7):70-77, July 2000.
- [2] A. Jalali, R. Padovani, R. Pankaj, "Data throughput of CDMA-HDR: A high efficiency high data rate personal communication wireless system," *Proceedings of the IEEE Vehicular Technology Conference*, Tokyo, Japan, May 2000.
- [3] "Physical Layer Aspects of UTRA High Speed Downlink Packet Access (Release 2000)", 3rd Generation Partnership Project Technical Specification Group, Radio Access Network, *3G Technical Report (TR) 25.848*, 2002.
- [4] G. L. Stuber, "Principles of Mobile Communications", Kluwer Academic Publishers, 1996.
- [5] M. Naghshineh and I. Katzela, "Channel assignment schemes for cellular mobile telecommunication systems: A comprehensive survey", *IEEE Personal Communications Magazine*, Vol. 3(3), pp 10-31, 1996.
- [6] "Ocelot network optimization tool", www.lucent.com.
- [7] "Network optimization tool", www.schema.com.
- [8] E. Amaldi et al, "Improved models and algorithms for UMTS radio planning", in *Proceedings of IEEE VTC*, Vol. 4, pp 2828-2832, Boston, Fall 2001.
- [9] I. Koutsopoulos and L. Tassulius, "Channel state-adaptive techniques for throughput enhancement in wireless broadband networks", in *Proceedings of INFOCOM*, 2001.
- [10] www.3gpp2.org, "CDMA2000 1X-DO standard", October 2001.
- [11] M Andrews et al., "Providing quality of service over a shared wireless link", *IEEE Communications Magazine*, 39(2):150-154, Feb 2001.
- [12] D.Tse, "Forward link multiuser diversity through rate adaptation and scheduling", *Bell Labs Presentation*, 1999.
- [13] H. Viswanathan and K. Kumaran, "Rate scheduling in multi-antenna packet data systems", submitted to *IEEE Transactions on Communications*, January 2002.
- [14] S. Shakkottai and A. Stolyar, "A study of scheduling algorithms for mixture of real and non-real time data in HDR", presented at the *17th International Teletraffic Congress*, September 2001.
- [15] Third Generation Partnership Project; Technical Specification Group Radio Access networks; "Evaluation Methods for High Speed Downlink Packet Access"; (3GPP document # R1-00-0909)

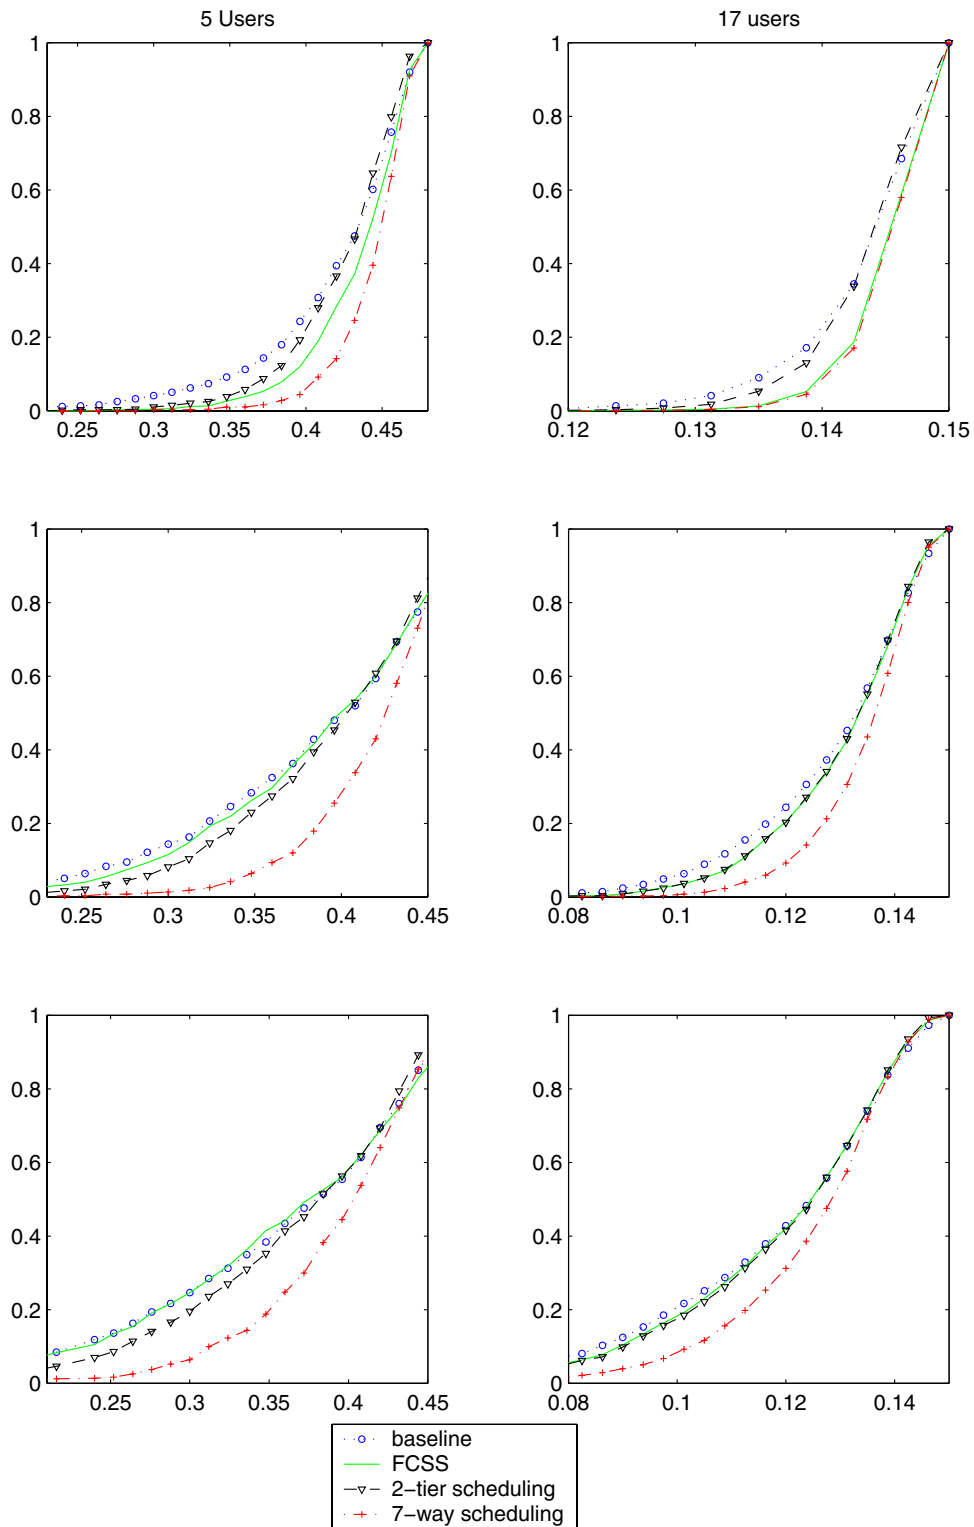


Fig. 3. Performance of various coordination schedules in a 7-cell configuration for uniform traffic density. The *x-axis* represents the per user throughput in bits/s/Hz and the *y-axis* represents the cumulative density function. Lognormal shadowing with variance of 4.3dB is assumed. The results are for 5 and 17 users and for each set of users Rayleigh fading and Rician Fading channels with K-factor 0dB and 6dB are simulated.

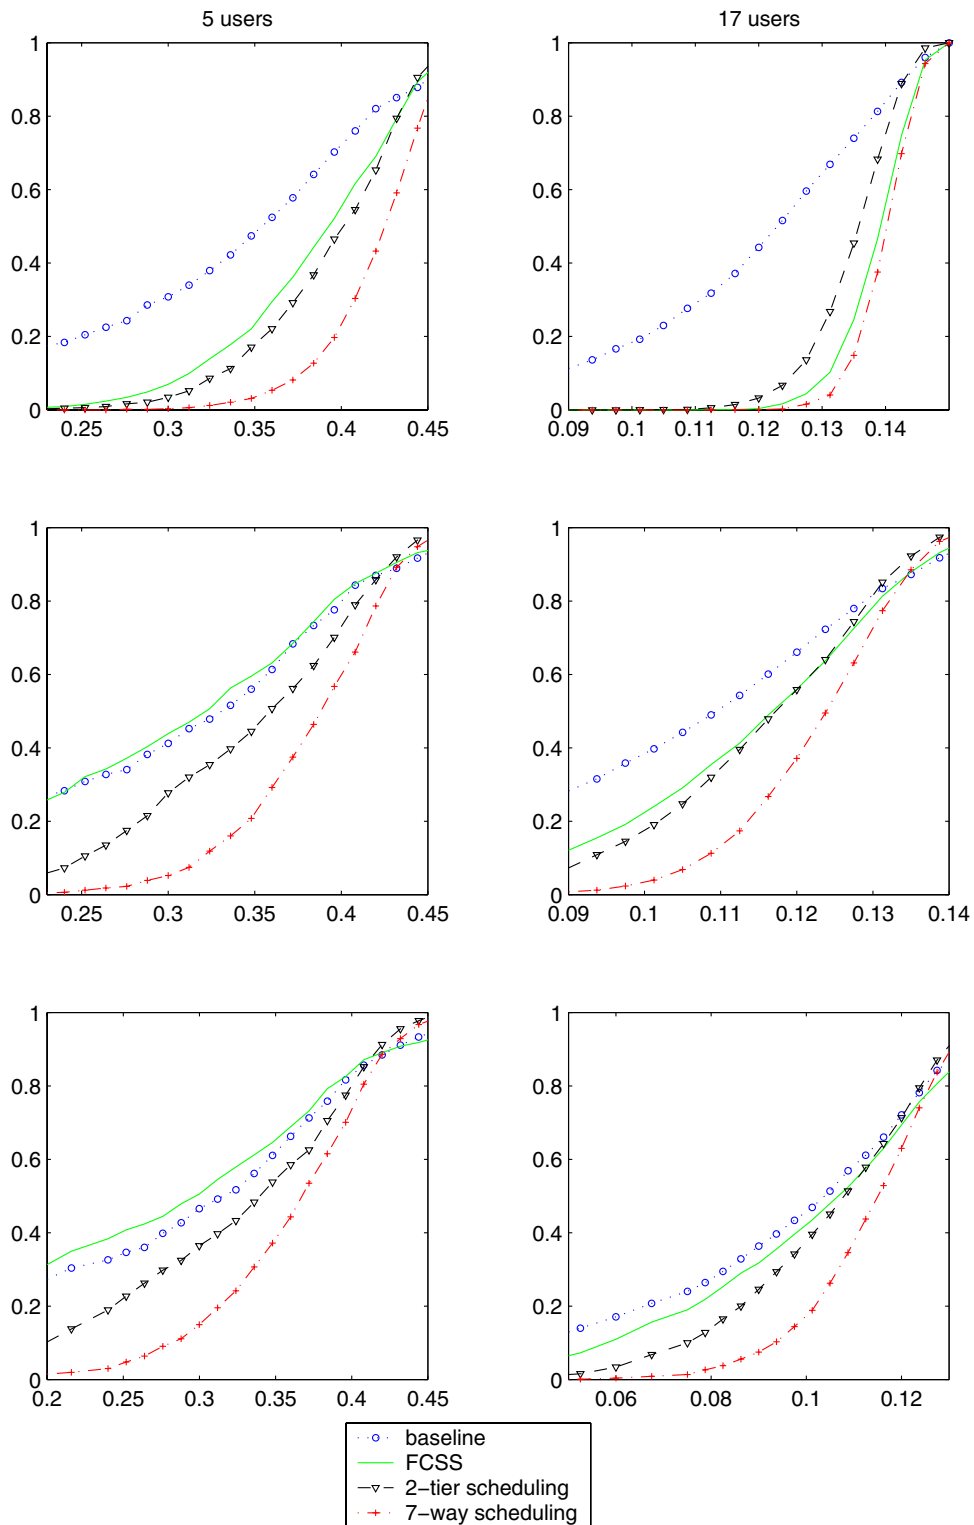


Fig. 4. Performance comparison of various scheduling algorithms with 7-cell coordination. Highly asymmetric traffic configuration is assumed, (99% of user in one cell). Cumulative distribution function of per user throughput with 5 and 17 cells per sector is simulated for Rayleigh and Rician fading channels with K-factor 0 and 6-dB.

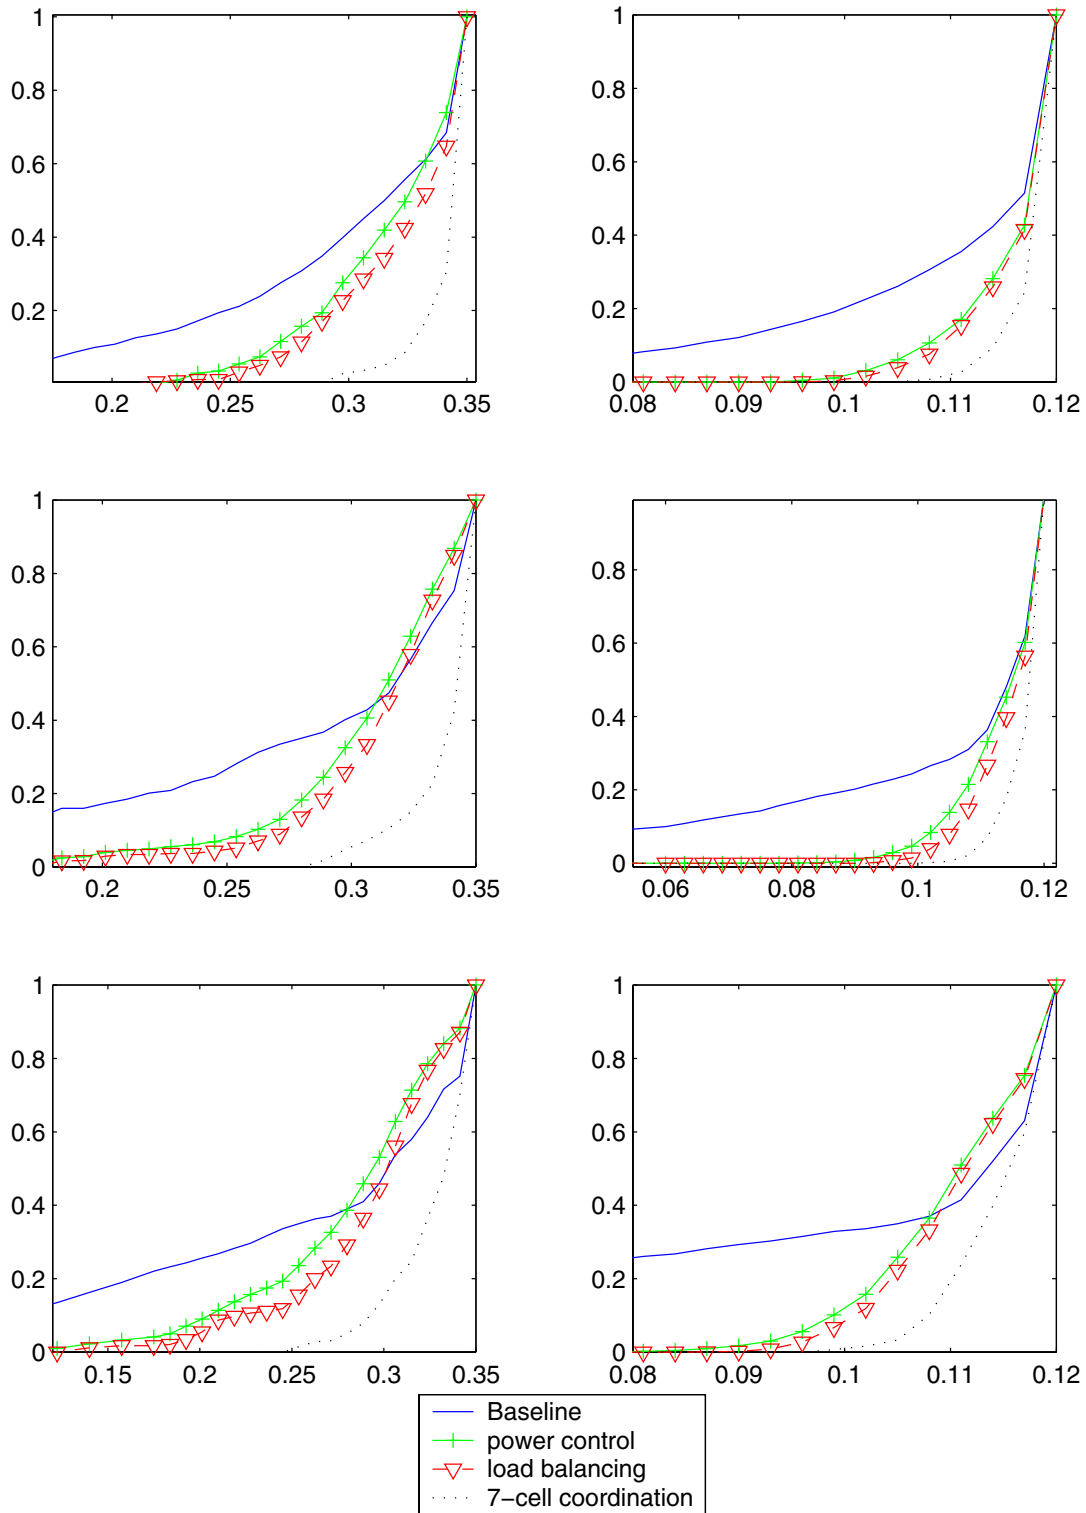


Fig. 5. Cumulative distribution function of per user throughput showing the effectiveness of load-balancing algorithm. Three different types of user densities are studied (uniform, and systems where 50% and 99% of total users are in the center cell). The plots on the left side are for 5 users per sector and the one on the right are for 17 users per sector.