Long-term Proportional Fairness Over Multiple Cells

Hatem Abou-zeid*, Stefan Valentin†, and Hossam S. Hassanein*

*Queen's University, Canada, h.abouzeid@queensu.ca, hossam@cs.queensu.ca
†Bell Labs, Alcatel-Lucent, Germany, stefan.valentin@alcatel-lucent.com

***Abstract*—In current cellular networks, Proportional Fair Scheduling (PFS) is performed separately at each base station. Such single-cell scheduling is not aware of the data rate that mobile users received in previously traversed cells. In this paper, we will show that, by lacking this information, Single-cell PFS fails to provide Proportional Fairness (PF) in the long-term as mobile users move over multiple cells. To overcome this limitation, we extend Single-cell PFS by the long-term average rate a mobile user received in previous cells. By being aware of this scheduling history, our Multi-cell PFS approach achieves substantially higher fairness and throughput compared to the traditional single-cell PFS approach.**

I. INTRODUCTION

Today's cellular networks are experiencing an explosion in data traffic, driven by large-screen devices (such as Smartphones, Tablets and Netbooks) and the growth of online media streaming [1]. Furthermore, network traffic is becoming more unevenly distributed in space and time [2] with high demand peaks moving across different cells at different times of the day.

To cope with the enormous traffic load, traditional macro cells are now accompanied by small cells [3] or new rollouts focus entirely on small-cell deployments. As a result, a user will traverse a larger number of cells per session. This increase in handovers per session is not only driven by the shrinking cell radius but also by the longer session time of modern handheld applications (e.g., video streaming, chatting). When users spend only a small fraction of their session time in each cell and move across a network with unbalanced load, this will lead to new challenges in providing long-term fair scheduling. This is because current schedulers operate at each cell independently.

Proportional Fairness (PF) is a common channel-state based scheduler employed at the Base Station (BS) for downlink scheduling. The scheduler exploits multiuser diversity to achieve a high throughput while maintaining fairness among the users. The intuition of the algorithm is to schedule users when they are at their peak achievable throughputs relative to their *averages*. This opportunistic approach results in a high network throughput. Moreover, if the users have statistically identical channels they will be served equally in the long run. However, a limitation of current networks is that PF scheduling is performed independently at each BS. When users with ongoing sessions move from one cell to another, the current cell is unaware of the average rate the users received in their previous cells, and therefore treats all incoming users equally. Users that have received a high average will compete for resources equally with users that have had low averages. As some users remain in congested cells for longer than others, the result is unfair scheduling when considering the long-term service over multiple cells. This problem becomes more significant in networks with unbalanced loads.

To overcome this limitation and provide long-term fairness, we present a simple coordination approach among multiple cells. Our *Multi-cell Proportional Fair Scheduling (MPFS)* strategy is aware of the average rate a user received in its previously traversed cells. By accounting for this *scheduling history*, MPFS can use capacity in a new cell to compensate a user for its poor performance in previous cells. As a result, the long-term fairness substantially increases for all users in the network. Moreover, as MPFS profits from the Channel Quality Information (CQI) statistics of all cells the user traverses, it can base its decisions on more accurate rate estimations. This even improves the average throughput of the network as we shall see in Sec. IV. Results show that the proposed scheduler can achieve a 20% throughput gain and a 30% fairness gain over current, single-cell PF scheduling. While reaching these substantial gains, this long-term multi-cell coordination comes at negligible computational effort and insignificant signalling overhead. We believe that ensuring long-term fairness over multiple cells will become increasingly relevant in future networks due to the current trends of smaller cell sizes, longer user sessions and higher user mobility.

We study our Multi-cell scheduler for the downlink and motivate it by the fact that traditional Single-cell Proportional Fair Scheduling (SPFS) fails to provide long-term Proportional Fairness over a network of mobile users. This violation of the PF property will be shown in Sec. IV-A. We, then, describe our MPFS approach in Sec. IV-B and take a closer look on its PF property and on the accuracy of its rate estimation in Sec. IV-C. The resulting performance is studied in Sec. V based on our system model from Sec. III. Before we detail the system model assumptions, let us first briefly review the related work.

This work was made possible by a *National Priorities Research Program* (NPRP) grant from the Qatar National Research Fund (Member of Qatar Foundation). This work is also supported in part by the program *RISE professional* of the DAAD and the program *PhD@BL* of Bell Labs, Stuttgart.

978-1-4673-2129-7/12/$31.00 ©2012 IEEE

II. RELATED WORK

Some research efforts have been made to coordinate the PF scheduling between multiple cells. These works however focus on instantaneous cooperation to achieve instantaneous objectives, i.e. BSs coordinate their transmissions at every transmission interval to minimize interference, balance load, or perform joint transmissions to a user such as in Coordinated Multi-Point (CoMP) [4].

In [5], Frank *et al.* propose a PF scheduling scheme for the 3GPP LTE uplink that takes inter-cell interference into account. Their scheme avoids high interference situations especially for cell-edge users, and therefore improves the average spectral efficiency. Bu *et al.* consider a network of BSs and propose a Load Balancing scheme that improves proportional fairness over the network by controlling the association of users among neighboring BSs [6]. Users will therefore not be associated to the BS that gives the strongest signal strength, but instead to the BS that satisfies a network-wide proportional fairness criterion. This scheme is extended in [7] where partial frequency reuse (an inter-cell interference mitigation mechanism) is jointly optimized with the load balancing in a multi-cell network. Therefore the authors consider both inter-cell association and intra-cell association in their solution. More recently, in [8] the authors consider the case where a user is served by multiple BSs simultaneously and extend PF scheduling to this architecture in order to provide an instantaneous proportional fairness over the network.

MPFS differs fundamentally from such multi-cell coordination approaches. Instead of adjusting scheduling at every transmission interval, our approach adjusts the scheduling when a user enters a new cell, based on the scheduling history in prior cells. The reason for this, is that we are interested in improving the long-term fairness between users as they traverse multiple cells as opposed to the instantaneous fairness at every transmission interval.

III. SYSTEM DESCRIPTION

Let us now briefly introduce our system assumptions, optimization objectives and metrics as well as the single-cell schedulers that refer as a baseline.

A. Models and Notation

We study a network with a Base Station set $\mathcal{M}$ and a user handheld set $\mathcal{N}$. An arbitrary user is denoted by $i \in \mathcal{N}$ and an arbitrary BS by $m \in \mathcal{M}$, where the number of BS is $|\mathcal{M}| = M$ and the number of users $|\mathcal{N}| = N$.

Each BS covers a hexagonal cell. All users move across the network according to the Random Way Point (RWP) mobility model with a constant speed S, zero pause time between the waypoints, and no wrap-around. Omitting the wrap-around creates a traffic hotspot in the center of the network, which allows us to study an uneven distribution of network load among the cells. Fig. 1 shows an example of a two-tier, 19 cell network with 3 user motion paths generated by the RWP model.

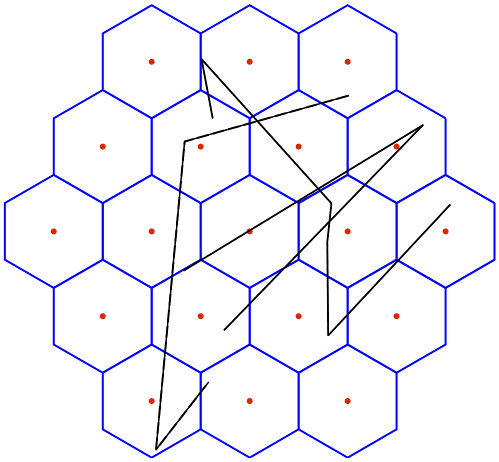

Fig. 1. The studied network with 19 hexagonal cells and the mobility paths of 3 exemplary users.

We model the wireless downlink as typical for studies on larger Long Term Evolution (LTE) systems. The path loss is calculated according to [9] as $\text{PL} = 128.1 + 37.6 \log_{10} d$ with the user-BS distance d in km. The BSs have omnidirectional antennas and a log-normal distribution with a variance of 8 dB accounts for slow fading. Fast fading is modeled as i.i.d. Rayleigh-fading, which implies fading is independent of the user speed. Note that such an assumption is feasible for systems such as LTE or WiMAX where interleaving, Hybrid Automatic Repeat Request (HARQ), or Space-Time Coding (STC) temporally de-correlate the channel coefficients over a Time Transmission Interval (TTI). Link adaptation is modeled using Shannon's equation where the SNR is clipped at 20 dB to account for a maximum modulation order of 64 Quadrature Amplitude Modulation (QAM).

For this initial study we assume a Full Buffer traffic model, meaning that each user has download traffic at any point in time.

B. Objectives and Metrics

In this paper, we study two network objectives (i) average network throughput maximization and (ii) network fairness maximization. The average network throughput T_{Net} is measured during the downlink as the sum of the average data rate taken over all users of the network. Fairness is measured as 5th percentile data rate over all users during the downlink, and is denoted as $T_{5\%}$. This fairness measure accounts for the throughput of users at the cell edge and is widely-used when studying cellular networks.

C. Single-cell Schedulers

We compare MPFS to SPFS to investigate the throughput and fairness gains that can be achieved by our extension of traditional SPFS. We also compare MPFS to Maximum Rate (MR) scheduling which serves as a reference for the maximum throughput achievable. All schedulers serve one user per TTI according to the users' CQI from the previous time slot.

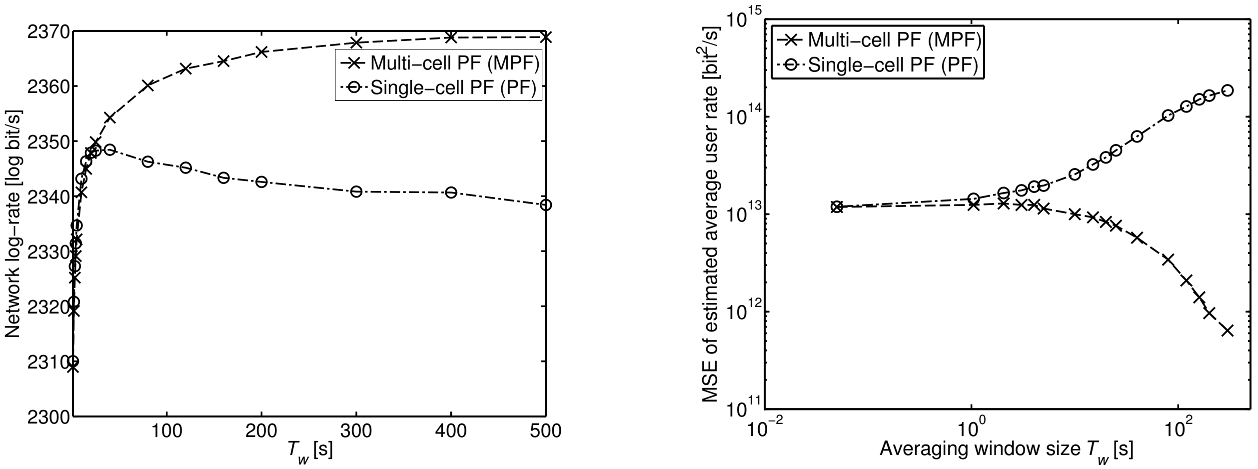

(a) Network logarithmic sum rate

(b) MSE of estimated average user rates

Fig. 2. Effect of the averaging window size T_w for MPFS and SPFS with $N = 250$ users, $S = 40\,\text{km/h}$.

In particular, MR schedules the user with highest instantaneous CQI, as observed from the previous TTI. This maximizes the sum throughput of the network but makes no effort to serve users fairly. Proportional Fairness (PF) aims for high throughput while maintaining fairness among the users. The intuition of the algorithm is to schedule users when they are at their peak rates relative to their average rates. In any TTI t, PF schedules the user $i^* = \arg\max_{\forall i \in \mathcal{N}} w_i(t)$ where the user weights are calculated $\forall i \in \mathcal{N} : w_i(t) = r_i(t)/\hat{R}_i(t)$. Here, $r_i(t)$ refers to the instantaneous data rate in the last time slot while

$$\hat{R}_i(t+1) = \frac{1}{T_w} r_i(t) p_i(t) + \left(1 - \frac{1}{T_w}\right) \hat{R}_i(t) \tag{1}$$

is the estimated moving average of the data rate. Here, $p_i(t)$ is a binary variable, which is equal to 1 when user i is scheduled at slot t and equal to 0 otherwise. Variable T_w denotes the observation time window, over which the moving average is computed.

It is important to note that the average rate $\hat{R}_i(t)$ is an *estimate* over T_w that may differ from the true average data rate R_i that the user experiences. We will elaborate on this difference in the following section.

IV. MULTI-CELL PROPORTIONAL FAIR SCHEDULING

We will now show that classic SPFS does not achieve proportional fairness over multiple cells. To overcome this limitation we extend SPFS to multiple cells. The resulting Multi-cell Proportional Fair Scheduling (MPFS) rule improves fairness and average throughput over the complete network for long session times.

A. Limitations of Single-cell PFS with Multiple Cells

In current networks, the average user rate $\hat{R}_i(t)$ in (1) is computed at each BS independently. Therefore, when a user with an ongoing session moves from one cell to another, $\hat{R}_i(t)$ from the previous cell is unknown to the current cell. This means that the current cell will restart the computation of (1) without the users past history, hence introducing errors in the estimation of $\hat{R}_i(t)$. As a result of this inaccurate estimation, SPFS cannot maintain Proportional Fairness (PF) over multiple cells and long-term fairness is lost.

To verify that SPFS cannot provide long-term PF over multiple cells, we test this scheduling rule against the PF property from [10, App. A]. A scheduler fulfills this property if it maximizes the logarithmic sum

$$R_{\log}^{\text{net}} = \sum_{i=1}^{N} \log R_i \tag{2}$$

of the average user rates

$$R_i = \frac{1}{T_w} \sum_{t=1}^{T_w} r_i(t) p_i(t) \tag{3}$$

if compared to all other schedulers for an asymptotically large T_w.

Fig. 2(a) shows our simulation results for increasing T_w. Here, the logarithmic sum rate $R_{\log}^{\text{net}}$ starts decreasing at $T_w = 20\,\text{s}$ when SPFS is employed. On the other hand, Multi-cell Proportional Fair Scheduling (MPFS) maintains a logarithmic shape and a substantially higher $R_{\log}^{\text{net}}$ than SPFS. As SPFS clearly does not maximize $R_{\log}^{\text{net}}$, it violates the PF property [10, App. A] when sum (2) is calculated over all users in the network. Thus, this observation proves that SPFS does not provide Proportional Fairness over multiple-cells.

SPFS fails to provide PF over the complete network since its average rate $\hat{R}_i(t)$ becomes inaccurate. As (1) only accounts for a single cell, each handover increases the error between $\hat{R}_i(t)$ and the true moving average $R_i(t)$. We measure this error as the Mean Square Error (MSE)

$$\text{MSE} := \frac{1}{T_w N} \sum_{t=1}^{T_w} \sum_{i=1}^{N} \left[\hat{R}_i(t) - R_i(t)\right]^2 \tag{4}$$

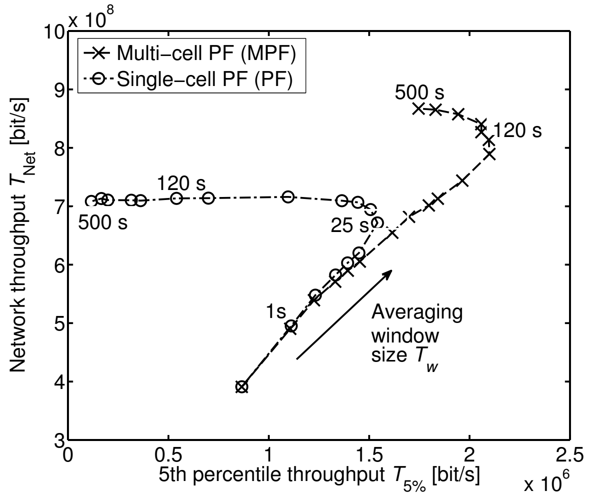

Fig. 3. Average network and 5th percentile throughput for increasing average window size T_w.

and plot it in Fig. 2(b). Since the probability that a user changes the cell increases for longer time windows, the MSE increases with T_w. Thus, the inaccurate moving average $\hat{R}_i(t)$ causes SPFS to violate the PF property (2). This is a clear motivation to revise the calculation of $\hat{R}_i(t)$ in (1).

Further, we make an initial investigation on how the average network throughput and fairness of SPFS changes with increasing T_w. Both metrics are shown in the scatter plot in Fig. **??**. Here, as expected the throughput initially increases with increasing T_w but then quickly saturates to 7×10^8 bit/s. More importantly, $T_{5\%}$ decreases significantly with increasing T_w, resulting in unfair scheduling. This is a direct consequence of the increase in the MSE when T_w increases as illustrated in Fig. 2(b).

B. Extending Single-cell to Multi-cell PF Scheduling

Let us now revise SPFS's calculation of the moving average (1) such that the users' scheduling history is included. To do so, we assume that either the user or the previous BS pass the latest $\hat{R}_i(t)$ to the new BS as soon as a user changes the cell. We illustrate this in Fig. 4 where a user changes from cell X to cell Y.

In this example, either BS_X or the user, reports the moving average $\hat{R}_{i,X}(t)$ to BS_Y. Here, we add index X to denote that $\hat{R}_{i,X}(t)$ was computed in cell X. Knowing $\hat{R}_{i,X}(t)$, BS_Y can now compute the scheduling weight for user i as $w_i = r_{i,Y}(t)/\hat{R}_{i,Y}(t)$ and the average rate as

$$\hat{R}_{i,Y}(t+1) = \frac{1}{T_w} r_{i,Y}(t) p_i(t) + \left(1 - \frac{1}{T_w}\right) \hat{R}_{i,X}(t) \quad (5)$$

where $\hat{R}_{i,Y}$ and $\hat{R}_{i,X}$ denote the average rates and $r_{i,Y}$ stands for the instantaneous rate during the downlink of user i in the respective cell.

Generally speaking, we change the weight calculation of SPFS to $\forall i \in \mathcal{N} : w_i(t) = r_{i,m(t)}(t)/\hat{R}_{i,m(t)}(t)$. Here, we include the scheduling history from the previously traversed cells by extending the moving average calculation of SPFS (1) to

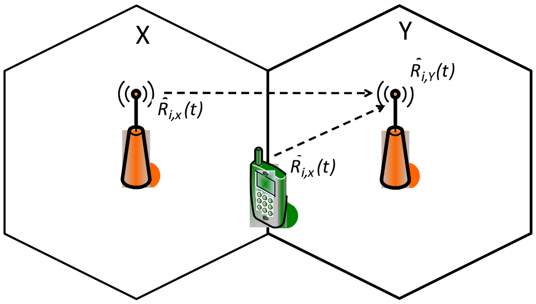

Fig. 4. As user i enters cell Y, $\hat{R}_{i,x}(t)$ is transmitted to BS_Y from either BS_X or the user terminal.

$$\hat{R}_{i,m(t)}(t+1) = \frac{1}{T_w} r_{i,m(t)}(t) p_i(t) + \left(1 - \frac{1}{T_w}\right) \hat{R}_{i,m(t-1)}(t) \quad (6)$$

where $m(t), m(t-1) \in \mathcal{M}$ are the respective BS indices in the current and previous time slot. If a user changes the cell, $m(t) \neq m(t-1)$. When a user does not change the cell $m(t) = m(t-1)$ and (6) reduces to (1).

From an implementation point of view, we have two options to perform the additional signalling of the average rates. Taking the example of Fig. 4, BS_X can signal the average rate $\hat{R}_{i,x}(t)$ to BS_Y via the core network, e.g., the X2 interface [11]. If such signalling is not supported, the user handheld can receive the average rate from BS_X and signal it to BS_Y during handover via the air link. Note that either of these options requires only an exchange of one average rate value per user, per handover. As a typical CQI resolution is 4 bits in LTE [12], MPFS adds only an insignificant signalling effort.

C. Effect of Multi-cell PF Scheduling

We now revisit Fig. 2 to verify if extending SPFS to multiple cells restores Proportional Fairness. Fig. 2(a) shows that Multi-cell Proportional Fair Scheduling (MPFS) maintains the logarithmic shape of (2) when T_w increases. Thus, Multi-cell PFS maximizes $R_{\text{log}}^{\text{net}}$ and, thus, fulfills the PF requirement when compared to SPFS.

The reason for restoring PF is shown in Fig. 2(b). Here, the MSE of MPFS is significantly smaller than with single-cell scheduling. By including the rates of the previous BSs, MPFS accurately computes the average rate. As the number of handovers per user increases in time, a larger T_w allows to collect more statistically independent rate samples from neighboring base stations. Consequently, the MSE of MPFS even decreases with larger observation windows and MPFS becomes more accurate with longer sessions.

The benefit on the average throughput and fairness is shown in Fig. 3. At small observation windows T_w, single and multi-cell scheduling achieve almost equal performance. Within this short time, handovers seldom occur and MPFS provides no benefit. This situation changes when T_w exceeds 20 s, which also represents the turnover point in Fig. 2(a) and 2(b).

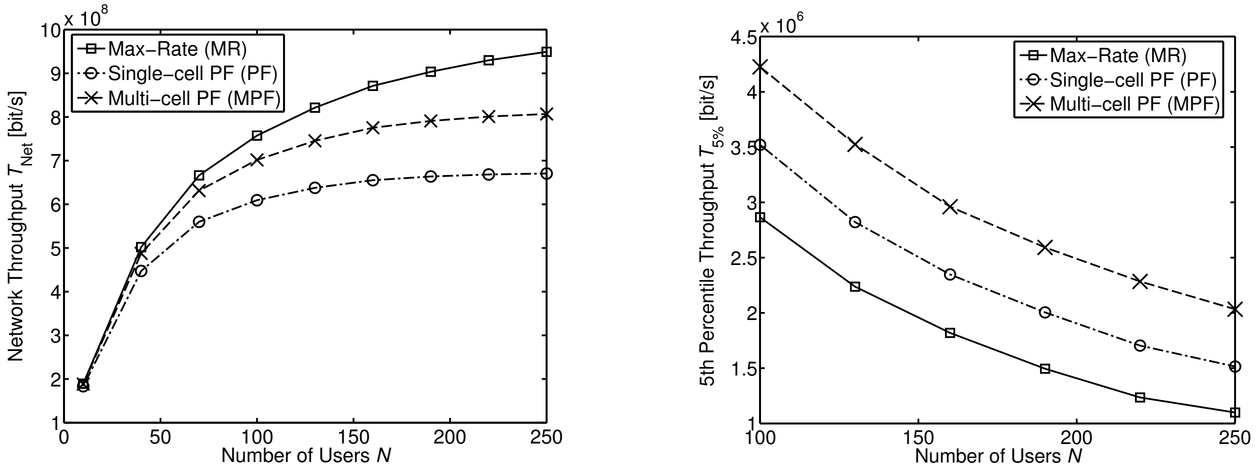

(a) Average network throughput T_{Net}

(b) 5th percentile throughput $T_{5\%}$

Fig. 5. Effect of the number of users N on the network throughput T_{Net} and 5th percentile throughput $T_{5\%}$ at a user speed of $S = 40\,\text{km/h}$. Here T_w is chosen such that $T_{5\%}$ is maximized.

As the users now frequently change among the cells, SPFS suffers from its inaccurate moving average. This results in incorrect scheduling decisions for SPFS, which is now clearly outperformed by MPFS in terms of average throughput and $T_{5\%}$ fairness. We will further study the performance of single and multi-cell scheduling in the following section.

V. PERFORMANCE EVALUATION

In this section we compare the average throughput T_{Net} and fairness $T_{5\%}$ of the proposed MPFS to SPFS and MR. We then investigate the effect user speed has on T_{Net} and $T_{5\%}$ for each of the schedulers. Finally, we present the potential energy savings of MPFS over SPFS.

A. Studied Factors and Parameters

We evaluate the schedulers in a network of 19 cells for varying network load and user speed. We parametrize the models from Sec. III by an inter-BS distance $D = 1\,\text{km}$, a BS transmit power of 40 W for the downlink, a center carrier frequency of 2 GHz, and a bandwidth of 10 MHz. Each simulation accounts for 900 s simulated time, which includes a 300 s warm-up period. Each simulation run is repeated 30 times with independent random numbers.

We compare MPFS to the single-cell schedulers MR and SPFS (cp. Sec. III). The averaging time window T_w is parametrized such that a particular scheduler achieves the optimization objective. That is, SPFS maximizes throughput at $T_w = 120\,\text{s}$ but requires $T_w = 25\,\text{s}$ to maximize the 5th percentile $T_{5\%}$ of the throughput as shown in Fig. 3. On the other hand, MPFS maximizes throughput at $T_w = 500\,\text{s}$ and $T_{5\%}$ at $T_w = 120\,\text{s}$. The values for T_w were obtained from previous simulations. As fairness is a long-term objective, $T_{5\%}$ is computed over 300 s.

B. Effect of Network Load

Fig. 5(a) and Fig. 5(b) show average throughput and fairness for varying load. The load can be entirely represented by the number of users N due to the Full Buffer traffic model. For both plots, T_w was chosen such that $T_{5\%}$ is maximized.

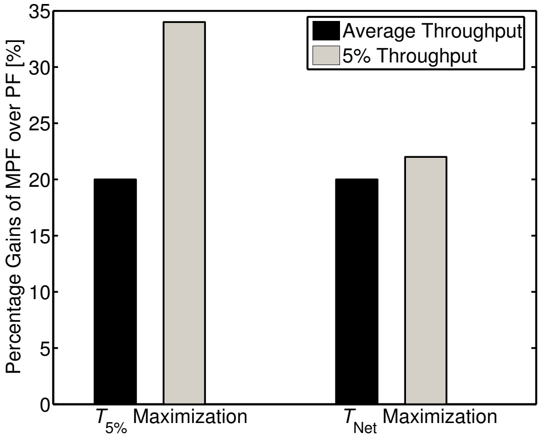

Fig. 6. The percentage gains of MPFS over SPFS. The gains are separately shown for the network objectives of $T_{5\%}$ maximization and T_{Net} maximization.

As shown in Fig. 5(a) MPFS clearly outperforms SPFS in terms of average network throughput T_{Net}. This gain exceeds 20% at 250 users. At low load, MPFS approaches the performance of MR. Our fairness results in Fig. 5(b) accompany these benefits. In terms of $T_{5\%}$, MPFS outperforms SPFS by 34% at high load. Thus, MPFS achieves higher throughput at the cell edge than all other studied single-cell schedulers. Very similar results were obtained for the case when T_w was chosen such that T_{Net} is maximized. A summary of the gains of MPFS over SPFS is presented in Fig. 6. We can see that the percentage gains in terms of average network throughput T_{Net} and $T_{5\%}$ are over 20% for both $\max T_{\text{Net}}$ and $\max T_{5\%}$ objectives (parametrization of T_w explained in Sec. V-A).

All in all MPFS achieves substantial average throughput and fairness gains when compared to SPFS. While the average

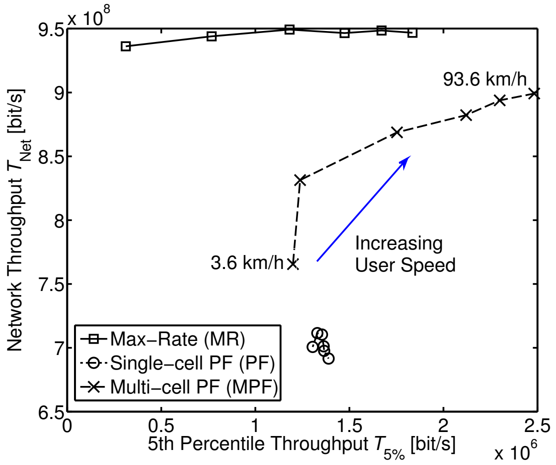

Fig. 7. Effect of user speed S on both the average network throughput T_{Net} and 5th percentile throughput $T_{5\%}$, for the different schedulers. $N = 250$.

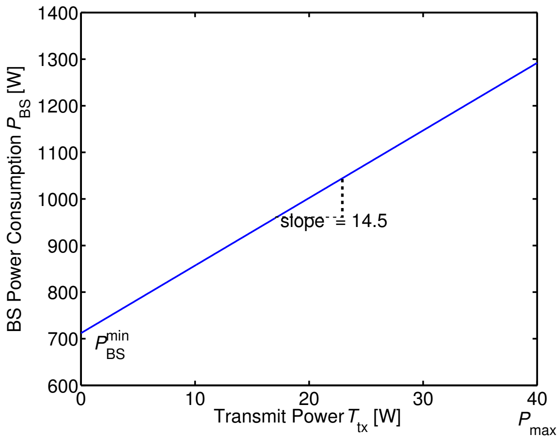

Fig. 8. Macro-cell BS power consumption P_{BS} for varying transmit power P_{tx} [13].

throughput gains increase with the number of users, MPFS improves fairness independent of the load. These gains result from an accurate estimation of the average rate by MPFS, which is aware of the users' rates in previous cells. These results confirm our expectations from Fig. 2.

C. *Effect of User Speed*

The scatter plot Fig. 7 shows how varying user speed affects T_{Net} and $T_{5\%}$ of the compared schedulers. A higher speed increases the number of handovers per session, which (i) reduces the CQI variance among the users by further distributing them in space, and (ii) improves the ergodicity of each user's channel by de-correlating its channel coefficient in time.

The de-correlated channel statistics increase the accuracy of the scheduler's rate estimation for a given number of CQI samples. Thus, with higher user speed, MPFS estimates $\hat{R}_{i,m(t)}$ more accurately and increases the 5th percentile throughput $T_{5\%}$. Even MR can profit from the reduced CQI variance (i), as shown by the fairness increase at higher speeds in Fig. 7. Such a gain cannot be found for SPFS. This scheme resets its scheduling history as soon as the user changes the cell and can, thus, not profit from the improved CQI statistics at higher user speed.

Fig. 7 also shows that T_{Net} increases significantly for MPFS. This is a result of the increased ergodicity of the channel (ii). This leads to additional multi-user diversity gains, which can be exploited by MPFS. MR can only slightly profit from (ii) as it already operates close to the ergodic capacity of the fading channel. SPFS provides no additional gains due to its erroneous estimate of $\hat{R}_i(t)$. By resetting this estimate at every handover, the estimation error increases at higher speeds when users change cells more frequently.

D. *Power Savings*

We now present the potential power savings of the Multi-cell Proportional Fair Scheduling (MPFS) scheduler compared to the SPFS. We use the BS power model proposed in [13] where the BS power consumption P_{BS} at a certain transmit power P_{tx} is determined by:

$$P_{\text{BS}} = P_{\text{BS}}^{\text{min}} + \Delta \cdot P_{\text{tx}}, \qquad 0 \leq P_{\text{tx}} \leq P_{\text{max}} \tag{7}$$

where P_{max} is the maximum transmit power, and $P_{\text{BS}}^{\text{min}}$ is the BS power consumption at the minimum possible transmit power as shown in Fig. 8. The Δ represents the slope of the linear relationship between the transmit power P_{tx} and BS power consumption P_{BS}. We obtain the values of $P_{\text{BS}}^{\text{min}}$ and Δ for a macro-BS from [13] and illustrate the resulting power model in Fig. 8. Also numerically, (7) becomes:

$$P_{\text{BS}} = 712 + 14.5 \cdot P_{\text{tx}}, \qquad 0 \leq P_{\text{tx}} \leq 40\,\text{W} \tag{8}$$

Fig. 9(a) shows how the network throughput varies with BS transmit power for the compared schedulers, and illustrates how we compute the power savings of the proposed Multi-cell Proportional Fair Scheduling (MPFS). The results of this figure are for a network objective of maximum throughput. To compute the power savings, we perform the following steps:

1) Determine the throughput of the SPFS scheduler at maximum transmit power (40 W), and then find the transmit power required by the MPFS to achieve an equal throughput, as indicated in Fig. 9(a).
2) Substitute the obtained values of $P_{\text{tx}}^{\text{SPFS}}$ and $P_{\text{tx}}^{\text{SPFS}}$ into (8), to compute the BS power consumption for each scheduler (which we denote by $P_{\text{BS}}^{\text{SPFS}}$ and $P_{\text{BS}}^{\text{MPFS}}$).
3) Calculate the power gain by: $(P_{\text{BS}}^{\text{SPFS}} - P_{\text{BS}}^{\text{MPFS}})/P_{\text{BS}}^{\text{SPFS}}$.

If on the other hand the network objective is to serve users fairly, then a similar procedure can be made on the results of the network for 5th percentile throughput maximization, as shown in Fig. 9(b). A summary of the percentage power savings for both network objectives is made in Fig. 10 where we can see that savings increase with increasing load and exceed 30% for 250 users. From Fig. 10 we can also observe that the gains when the network objective is to maximize

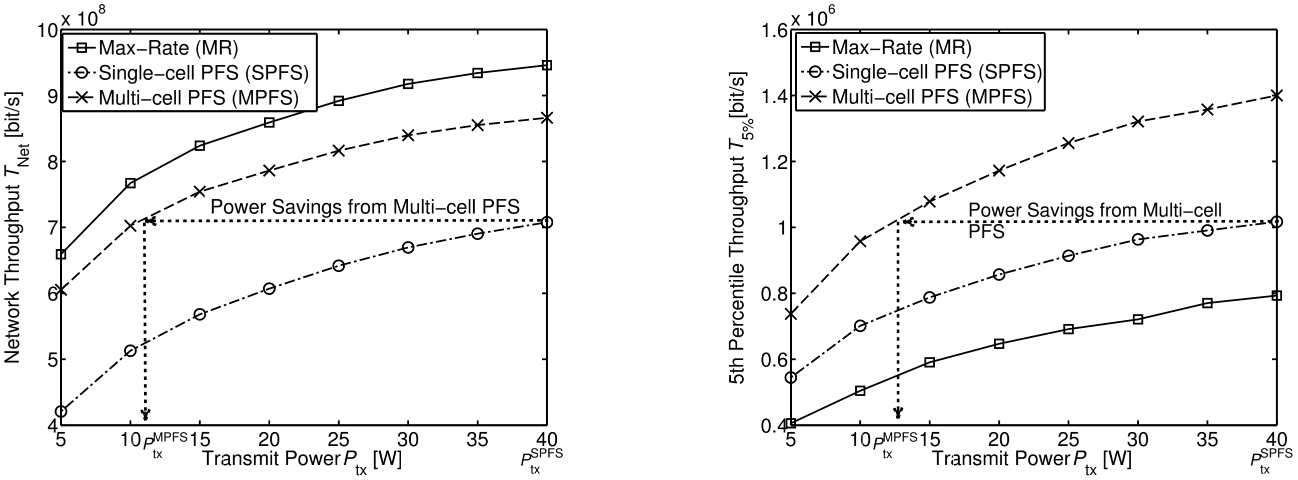

(a) Average network throughput T_{Net}

(b) 5th percentile throughput $T_{5\%}$

Fig. 9. Effect of the BS transmit powers P_{tx} on the network throughput T_{Net} and 5th percentile throughput $T_{5\%}$ at a load of $N = 250$ users. Here T_w is chosen such that T_{Net} is maximized.

throughput are slightly higher than the case for maximizing 5th percentile throughput.

VI. CONCLUSION

Although classic wireless schedulers provide Proportional Fairness (PF) within a single cell, they fail to keep this fairness property when users move among multiple cells. To overcome this problem, we extend classic Single-cell Proportional Fair Scheduling (SPFS) by the data rates from previously visited cells. The resulting Multi-cell Proportional Fair Scheduling (MPFS) strategy is aware of a user's scheduling history and improves long-term fairness over the complete network.

By signalling only one additional rate value per user and handover, our extension is feasible and provides significant gains. Compared to the traditional single-cell approach, MPFS increases the data rate by up to 20% on the average and by up to 34% at the cell edge. These gains are achieved simultaneously and increase with the user speed, traffic load, and session time. We expect even higher performance gains with upcoming small-cell deployments [3], where the average number of handovers per session will further increase.

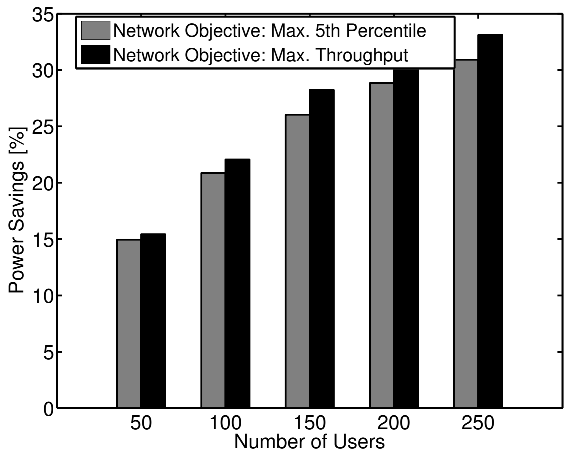

Fig. 10. MPFS power savings for different number of users.

REFERENCES

[1] G. Maier, F. Schneider, and A. Feldmann, "A first look at mobile handheld device traffic," in *Proceedings of the 11th International Conference on Passive and Active Network Measurement (PAM '10)*, vol. 6032, pp. 161–170, Apr. 2010.

[2] U. Paul, A. P. Subramanian, M. M. Buddhikot, and S. R. Das, "Understanding traffic dynamics in cellular data networks," May 2011.

[3] A. Damnjanovic, J. Montojo, Y. Wei, T. Ji, T. Luo, M. Vajapeyam, T. Yoo, O. Song, and D. Malladi, "A survey on 3GPP heterogeneous networks," *IEEE Wireless Communications*, vol. 18, pp. 10–21, June 2011.

[4] P. Marsch and G. Fettweis, eds., *Coordinated Multi-Point in Mobile Communications: From Theory to Practice*. Cambridge University Press, 2011.

[5] P. Frank, A. Muller, H. Droste, and J. Speidel, "Cooperative interference-aware joint scheduling for the 3gpp lte uplink," in *Personal Indoor and Mobile Radio Communications (PIMRC), 2010 IEEE 21st International Symposium on*, pp. 2216 –2221, sept. 2010.

[6] T. Bu, L. Li, and R. Ramjee, "Generalized proportional fair scheduling in third generation wireless data networks," in *Proc. IEEE INFOCOM*, April 2006.

[7] K. Son, S. Chong, and G. Veciana, "Dynamic association for load balancing and interference avoidance in multi-cell networks," *IEEE Transactions on Wire*, vol. 8, pp. 3566–3576, 2009.

[8] H. Zhou, P. Fan, and J. Li, "Global proportional fair scheduling for networks with multiple base stations," *IEEE Transactions on Vehicular Technology*, vol. 60, pp. 1267–1879, 2011.

[9] 3GPP, "Physical layer aspect for E-UTRA," Technical Report TR 25.814 V7.1.0, 3GPP, Oct. 2006.

[10] P. Viswanath, D. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Transactions on Information Theory*, vol. 48, pp. 1277–1294, June 2002.

[11] E. Dahlman, S. Parkvall, J. Skold, and P. Beming, *3G Evolution: HSPA and LTE for Mobile Broadband*. Academic Press, 2008.

[12] 3GPP, "E-UTRA physical layer procedures (release 10)," Technical Specification TS 36.213 V10.4.0V, 3GPP, Dec. 2009.

[13] EARTH, "Energy efficiency analysis of the reference systems, areas of improvements and target breakdown," Technical Report INFSO-ICT-247733, Deliverable D2.3, EARTH, Jan. 2012.