

Optimal Predictive Resource Allocation: Exploiting Mobility Patterns and Radio Maps

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

*Queen’s University, Canada, h.abouzeid@queensu.ca, hossam@cs.queensu.ca

†Bell Labs, Alcatel-Lucent, Germany, stefan.valentin@alcatel-lucent.com

Abstract—Resource Allocation (RA) in cellular networks is a challenging problem due to the demanding user requirements and limited network resources. Moreover, mobility results in channel gains that vary significantly with time. However, since location and received signal strength are correlated, user mobility patterns can be exploited to predict the data rates they will experience in the future. In this paper, we show that with such predictions, long-term RA plans that span multiple cells can be made. We formulate an optimal Predictive Resource Allocation (PRA) framework for a network of cells as a linear programming problem for three different objectives. Presented numerical results provide a benchmark of the PRA performance in realistic and random user mobility scenarios. Significant network and user satisfaction gains are observed compared to RA schemes that do not utilize any predictions.

I. INTRODUCTION

Coping with the unprecedented growth of mobile traffic is a predominant challenge in cellular networks. Furthermore, vehicular users experience strong data rate fluctuations as their locations change rapidly. However, a correlation exists between geographic location and the received signal strength. This has facilitated the creation of *radio maps* such as [1] where average values of historic signal strengths are stored on a map. Studies have also shown that peoples’ daily routes exhibit a high degree of temporal and spatial regularity, with people inclined to follow particular routes to and from frequently visited places [2]. Therefore, if these mobility patterns are coupled with radio maps, the average data rates users will experience along a trip can also be predicted.

In cellular networks, the goal of Resource Allocation (RA) is to distribute Base Station (BS) bandwidth and power efficiently and fairly while satisfying individual user Quality of Service (QoS) requirements. Extensive work on RA has been presented in literature with different objectives and algorithm complexities [3], [4]. Whereas existing works focus on achieving instantaneous objectives, in this paper, we exploit mobility patterns and radio maps to make long-term Predictive Resource Allocation (PRA) plans. With such information, BSs can know if users are headed to coverage holes, or to network congested zones. Fig. 1 illustrates the relative signal strength along a road network, and shows two use-cases for PRA. In case A, BSs can prioritize users headed to low rate areas (marked in dark colors) and pre-allocate the required data before hand. This is particularly useful for buffered media streams to ensure smooth playback when users enter the low data rate areas. In case B, BSs can plan resource reservations




Fig. 1. Use of predictive resource allocation.

for users arriving from low rate areas, for instance to accelerate the progress of a file download that was stalled. The goal is to improve user satisfaction by making long-term RA provisions.

Towards this end, we present an optimal PRA framework that: 1) extends the RA planning horizon to tens of seconds, that span multiple cells, 2) exploits predictions of users’ future data rates to make long-term RA at each BS, and 3) introduces long-term multi-cell cooperation, where allocations made to users in one cell impact the amounts allocated in future cells traversed by the users. The framework is formulated as a linear programming problem for three different network objectives. The first objective maximizes the throughput of multiple cells for a specified time horizon, and the second achieves long-term max-min fairness, also over multiple cells. The third objective minimizes degradations of buffered video streams by performing the optimal pre-allocations to prevent video stalling. In this work we provide the upper bounds of the multi-cell PRA and therefore assume ideal predictions of user rates. We present performance benchmarks of each PRA objective for two mobility scenarios: a road network with realistic vehicle mobility, and a 19 cell network with Random Way Point (RWP) mobility.

In the next section we review related work, and follow that with our system models in Sec. III. In Sec. IV we formulate the multi-cell PRA problem, and discuss the numerical results in Sec. V. Finally, we conclude in Sec. VI.

II. RELATED WORK

Multi-cell resource allocation is an active research topic which has focused primarily on instantaneous cooperation

to achieve short-term objectives, e.g. BSs coordinate their transmissions periodically to minimize interference, balance load, or perform joint transmissions to a user [5]. This form of multi-cell RA differs from that discussed in this paper. Our goal is to exploit user rate prediction to make optimal long-term allocations, e.g. prioritizing users entering coverage holes or congested cells. In a related work, we present multi-cell RA that exploits user QoS history in previous cells, to make allocations in the current cell [6]. Such a scheme improves long-term QoS by prioritizing users that have had poor service in previously traversed cells. However, it does not exploit rate predictions to make long-term plans.


User mobility and rate *predictions* have been used in a limited number of recent works to improve network performance. For example, Ali et al. use rate predictions to increase system data rate of a single-cell in [7]. This work does not consider cooperation between multiple cells, or formulate different network objectives, and assumes an infinite backlog traffic model only. In [8], Yao et al. demonstrate the use of radio maps in video streaming where video servers proactively switch to the transmission rates predicted to be supportable at the user, and therefore improve TCP rate control and throughput. However, they do not investigate the multi-user RA problem, i.e. how resources are distributed among users with time to improve QoS. In a recent work, we presented an introductory proposal of exploiting user mobility predictions for multi-cell RA in [9].

III. SYSTEM DESCRIPTION

A. System Model

We consider a network with a BS set \mathcal{M} of M BSs and a user set \mathcal{N} with N users. An arbitrary BS is denoted by $j \in \mathcal{M}$ and a user by $i \in \mathcal{N}$. Users are associated to BSs based on the strongest received signal. The set $\mathcal{U}_{j,t}$ contains the indices of all the users associated with BS j at time t .

Two network and mobility scenarios are considered. The first is a network of six BSs shown in Fig. 2(a), that covers the illustrated road network. To provide realistic vehicular mobility we use the SUMO traffic simulator [10] to generate traces for the three routes denoted by A, B and C. Vehicles



(a) Road network with mobility routes.

traverse these routes with equal probability. To provide system evaluation in a more general network, we also model the 19 cell network illustrated in Fig. 2(b). Users move according to the RWP mobility model with a constant speed S , zero pause time between the waypoints, and no wrap-around. While this model is not practical, it enables the study of PRA when users experience independent sequences of data rate fluctuations.

A radio map $\text{RMap}_{x,y}$ is created with the receive data rates at all the (x,y) coordinates that users can traverse using the path loss model $\text{PL} = 128.1 + 37.6 \log_{10} d$ where the BS-user distance d is in km [11]. The data rate is computed using Shannon's equation with SNR clipping at 20 dB to account for practical modulation orders. Therefore, a user i at location (x,y) at time t , will have a feasible data rate $\hat{r}_{i,t} = \text{RMap}_{x,y}$. In this study, we demonstrate the potential gains of the optimal scenario. Therefore, we assume an accurate radio map and predictable user routes (implying predictable rates) for a duration of T seconds, which we call the *prediction window*.


Time t is divided in equal slots of duration Δt , during which $\hat{r}_{i,t}$ can be assumed to be constant. We set $\Delta t = 1$ s as vehicle speeds can reach 20 m/s and therefore path loss can be affected at a granularity of one second. During Δt , the BS channel can be shared in arbitrary ratios between the users. We define the resource sharing factor $x_{i,t+\Delta t} \in [0,1]$ as the fraction of time (during the slot Δt) that the BS bandwidth is assigned to user i . This is the optimization variable used by PRA to define long-term user allocations over multiple BSs. Note that using a similar PRA framework, other resource sharing factors pertaining to specific radio access technologies can also be studied.

B. Reference Resource Allocators

To provide a performance reference we consider the following resource allocators.

1) *Max-rate Allocation*: In Maximum Rate (MR) allocation, the user with the highest data rate $\hat{r}_{i,t}$ is granted full channel access, i.e. $x_{i,t} = 1$. This maximizes the network throughput but makes no effort to serve users fairly.

2) *Equal Share Allocation*: In Equal Share (ES) allocation, the BS resource is shared equally among the users at each time



(b) The 19 cell network with RWP mobility.

Fig. 2. Network and mobility models considered.

slot t . If there are $N_{j,t}$ users associated with BS j at time t (i.e. $N_{j,t} = |\mathcal{U}_{j,t}|$), then $x_{i,t} = 1/N_{j,t}$ for each user $i \in \mathcal{U}_{j,t}$, and the rate received by each is $\hat{r}_{i,t}/N_{j,t}$.

3) *Rate-Proportional Allocation*: We propose the Rate-Proportional (RP) allocator to achieve high throughput while still serving users with poor channel conditions. The resource assigned to each user i at time t is in proportion to the achievable data-rate of that user $\hat{r}_{i,t}$. Therefore, $x_{i,t} = \hat{r}_{i,t} / \sum_{i \in \mathcal{U}_{j,t}} \hat{r}_{i,t}$, and the user rate received is $x_{i,t} \hat{r}_{i,t}$.

IV. MULTI-CELL PREDICTIVE RESOURCE ALLOCATION: PROBLEM FORMULATION

Predictive Resource Allocation exploits predictions of user movement and a road radio map, to optimize long-term RA as users traverse a region covered by multiple BSs. These predictions are assumed to be available at a central coordinating BS which plans user allocations for all the BSs in the region of interest. The objective of PRA is to optimize user allocations based on both, their long-term data requirements, and whether they are headed to better or worse channel, and network congestion conditions. In this paper, we study the gains that can be achieved by exploiting this information, and present three PRA formulations with different objectives.

A. Preliminaries and Assumptions

The data traffic requested by user i at time t is denoted by $D_{i,t}$, and the total data requested during T is denoted by D_i . We consider three cases for user traffic: (i) full buffer traffic, where $D_{i,t} \rightarrow \infty \forall i, t$, (ii) file download traffic, where D_i is finite $\forall i$, and (iii) buffered video streams, where $D_{i,t} = V \forall i, t$, where V denotes the target streaming rate.

In this work we are interested in studying the PRA schemes under ideal conditions. Therefore, the predicted user rates $\hat{r}_{i,t}$ are assumed to be accurate. The output of the PRA, the resource sharing factor $x_{i,t} \in [0, 1]$, is determined for the prediction window T , for all users. Fig. 3 illustrates an example of $x_{i,t}$ for a network with two BSs and four users. Users are assigned resources depending on the network objectives and their application needs. A summary of the commonly used notation in this paper is provided in Tab. I.

B. PRA-MaxNetRate

The objective of PRA-MaxNetRate is to maximize the total network throughput over the duration T . At time t , for the prediction interval $[t+1, t+T]$, it can be formulated as the Linear Program (LP):

$$\begin{aligned} & \underset{x_{i,t+\Delta t}}{\text{maximize}} && \sum_{\Delta t=1}^T \sum_{i=1}^N r_{i,t+\Delta t} x_{i,t+\Delta t} && (1) \\ & \text{subject to:} && \text{C1: } \sum_{i \in \mathcal{U}_{j,t}} x_{i,t+\Delta t} \leq 1, && \forall j, \Delta t \\ & && \text{C2: } \sum_{\Delta t=1}^T r_{i,t+\Delta t} x_{i,t+\Delta t} \leq D_i, && \forall i \\ & && \text{C3: } x_{i,t+\Delta t} \hat{r}_{i,t+\Delta t} \geq \text{MBR}, && \forall i, \Delta t \\ & && \text{C4: } 0 \leq x_{i,t+\Delta t} \leq 1 && \forall i, \Delta t. \end{aligned}$$




Fig. 3. Example of the resource sharing factor $x_{i,t}$ in PRA.

Note that the outer summation over time in the objective is included as PRA optimization is made for all time slots in the prediction window T . C1 expresses the resource sharing limitation at each BS. It ensures that the sum of the resource sharing factors of all users associated with a BS j is equal to 1 at every time slot. As C1 is applied at each BS, and every time slot, there are MT constraints in total. C2 limits the amount of data assigned to each user during T to the total amount request. This leads to N constraints. In C3, a Minimum Bit Rate (MBR) is defined to provide a lower limit of user allocation for each time slot. If this is set to zero, we get an upper bound of the PRA-MaxNetRate objective. C4 provides the bounds for the resource sharing factor. In total there are $MT + N + 3NT$ constraints and $2NT$ variable bounds. Note that the objective function couples rates received by users from several BSs during T , and is solved centrally for all cooperating BSs.

In PRA-MaxNetRate, a user in a congested, low-rate region that is heading to a low-density, high-rate region, will only be served with the MBR until the better conditions commence. This to ensure maximum network utilization when the user arrives at the high-rate region. PRA-MaxNetRate provides the upper bound on the throughput of any formulation in the PRA framework. However, if all the users have full buffer traffic, the throughput will be equal to the traditional MR allocation discussed in Sec. III-B.

C. PRA-MaxMin Fairness

PRA-MaxMin has the objective of achieving long-term max-min fairness among the users as they traverse multiple BSs with different congestion and road signal strengths. It allocates resources such that the minimum user throughput during T is maximized. This is formulated as:

$$\begin{aligned} & \underset{x_{i,t+\Delta t}}{\text{maximize}} && \min_i \frac{1}{D_i} \sum_{\Delta t=1}^T r_{i,t+\Delta t} x_{i,t+\Delta t} && (2) \\ & \text{subject to:} && \text{C1, C2, C3, C4,} \end{aligned}$$

where the constraints C1-C4 are similar to the PRA-MaxNetRate constraints in (1). Note that the objective in (2) achieves a max-min allocation of the *sum* of the data allocated

TABLE I
SOME NOTATIONS USED IN THIS PAPER

Symbol	Description
T	Prediction window: time duration over which user routes are predictable [s]
t	Current time slot: each slot has a duration of 1 s
$\mathcal{U}_{j,t}$	Set of the indices of users associated with BS _{j} at time t
$\hat{r}_{i,t}$	Predicted rate that user i can receive at time t [bits/s]
$x_{i,t}$	Fraction of BS resource assigned to user i at time t
V	Video streaming rate [bits/s]
$D_{i,t}$	Traffic requested by user i at time t [bits]
D_i	Total traffic requested by user i during T s [bits]
$\tilde{D}_{i,t}$	Cumulative (video) traffic requested by user i at time t [bits]
$VD_{i,t}$	Video degradation perceived by user i at time t

to each user during T , and not the individual data at each time slot t . This achieves the desired long-term fairness, while C3 provides the MBR at each time slot. The variable D_i in the optimization objective is used to normalize the allocated user data by the total data requested during T . We are interested in maximizing the minimum of this ratio over all users. As PRA-MaxMin has a piece-wise linear concave objective function, it can be expressed as an equivalent LP by introducing an optimization variable Y and constraint C5, as follows:

$$\underset{x_{i,t+\Delta t}}{\text{maximize}} \quad Y \quad (3)$$

subject to: C1, C2, C3, C4

$$\forall i \quad \text{C5:} \quad -\frac{1}{D_i} \sum_{\Delta t=1}^T r_{i,t+\Delta t} x_{i,t+\Delta t} + Y \leq 0.$$

PRA-MaxMin allocation can be used for best effort, delay tolerant applications such as FTP downloads or software updates, where MBR may even be zero. The network can then schedule long-term predictive allocations to ensure that users are served equally as they move across multiple cells.

Other forms of long-term fairness can also be achieved with PRA by modifying the objective and/or constraints. For example, a log utility can be used in the objective of (1) to achieve a notion of predictive-proportional fairness.

D. PRA-Video Degradation Minimization

Definitions: The cumulative data requested by a user at each time slot is defined as $\tilde{D}_{i,t+\Delta t} = \sum_{t'=t+1}^{t+\Delta t} D_{i,t'}$, where $\Delta t = 1, 2, \dots, T$ and $D_{i,t'} = V$ for video streaming. This is illustrated in Fig. 4 as the minimum target cumulative data for smooth video playback. Similarly, the cumulative data received, given a user resource allocation $x_{i,t+\Delta t}$, is defined as $\tilde{R}_{i,t+\Delta t} = \sum_{t'=t+1}^{t+\Delta t} x_{i,t'} \hat{r}_{i,t'}$. If this is higher than $\tilde{D}_{i,t+\Delta t} \forall \Delta t$ then the video is played smoothly, and the difference $\tilde{R}_{i,t+\Delta t} - \tilde{D}_{i,t+\Delta t}$ represents the amount of video content that is pre-buffered at the end user at time $t + \Delta t$. However, if $\tilde{R}_{i,t+\Delta t} < \tilde{D}_{i,t+\Delta t}$, the user experiences video

stalling, or a lower quality video, and therefore the video experience is degraded. We define Video Degradation (VD) as the amount of unfulfilled video demand. For a given predicted rate $\hat{r}_{i,t+\Delta t}$, and user allocation $x_{i,t+\Delta t}$, it is defined as:

$$VD_{i,t+\Delta t} = [\tilde{D}_{i,t+\Delta t} - \tilde{R}_{i,t+\Delta t}]^+ \quad (4a)$$

$$= \left[\sum_{t'=t+1}^{t+\Delta t} D_{i,t'} - \sum_{t'=t+1}^{t+\Delta t} x_{i,t'} \hat{r}_{i,t'} \right]^+ \quad (4b)$$

where $[x]^+ = \max\{0, x\}$. Fig. 4 shows a sample user allocation, where degradation occurs between 150 s and 220 s.

Optimization Problem: The objective of this PRA scheme is to make the optimal pre-buffering allocations to users, in advance, so that they all experience smooth playback. If this is not possible (e.g. at high load), then the total amount of video degradation experienced by the network is minimized. The optimization problem, which we refer to as PRA-MinVD, is formulated as:

$$\underset{x_{i,t+\Delta t}}{\text{minimize}} \quad \sum_{\Delta t=1}^T \sum_{i=1}^N \left[\sum_{t'=t+1}^{t+\Delta t} D_{i,t'} - \sum_{t'=t+1}^{t+\Delta t} x_{i,t'} \hat{r}_{i,t'} \right]^+ \quad (5)$$

subject to: C1, C2, C4.

The objective function is piece-wise linear and convex, and can therefore be reformulated as an LP as follows:

$$\underset{x_{i,t+\Delta t}, \text{Deg}_{i,t+\Delta t}}{\text{minimize}} \quad \sum_{\Delta t=1}^T \sum_{i=1}^N \text{Deg}_{i,t+\Delta t} \quad (6)$$

subject to: C1, C2, C4

$$\forall i, \Delta t \quad \text{C5:} \quad \tilde{D}_{i,t+\Delta t} - \sum_{t'=t+1}^{t+\Delta t} x_{i,t'} \hat{r}_{i,t'} - \text{Deg}_{i,t+\Delta t} \leq 0$$

$$\forall i, \Delta t \quad \text{C6:} \quad \text{Deg}_{i,t+\Delta t} \geq 0.$$

Here we have introduced degradation $\text{Deg}_{i,t+\Delta t}$ as additional optimization variables which we restrict to have positive values in C6. The objective now, is to minimize the sum of $\text{Deg}_{i,t+\Delta t}$, which is linear. This is equivalent to the objective in (5) as constraint C5 dictates that $\text{Deg}_{i,t+\Delta t}$ is positive only if




Fig. 4. Accumulative user data allocation in video streaming.

$\tilde{D}_{i,t+\Delta t} > \tilde{R}_{i,t+\Delta t}$, which represents the amount of VD defined in (4a). On the other hand, if $\tilde{D}_{i,t+\Delta t} < \tilde{R}_{i,t+\Delta t}$, C5 will be satisfied for any value of $\text{Deg}_{i,t+\Delta t} \leq \tilde{R}_{i,t+\Delta t} - \tilde{D}_{i,t+\Delta t}$. In this case, $\text{Deg}_{i,t+\Delta t}$ will be set to zero (indicating no degradation), as the objective is to minimize $\text{Deg}_{i,t+\Delta t}$, and $\text{Deg}_{i,t+\Delta t}$ cannot be negative by C6. Note that C3 is not included since buffered video streams have their cumulative demand expressed in C5, and therefore the per time slot MBR in C3 is not needed. Note that extensions to (6) that account for other aspects of perceived video QoS are also possible.

Thus, PRA can be used to improve network throughput, long-term user fairness, and QoS for buffered video streams. PRA formulations are also possible for other QoS objectives.

V. SIMULATIONS AND NUMERICAL RESULTS

In this section, we use the presented optimization framework as a benchmarking tool to study the performance upper bound of PRA for different network objectives and densities.

A. Simulation Set-up

We evaluate PRA for both networks of Fig. 2 where the inter-BS distance is set to 1 km for the 19 cell network. BS transmit power is 40W and the bandwidth is 10 MHz. In the RWP scenario, user speed S is 10 m/s, file download size is 1 Gbits, and the streaming rate V is 4 Mbps. For the road network, user speed is variable as obtained from the SUMO trace file, file sizes are 750 Mbits and V is 3 Mbps. The prediction window T is 250 s for both networks. PRA-MaxNetRate and PRA-MaxMin have an MBR of zero to illustrate their performance bound.

B. Performance Metrics

- T_{Net} : the average downlink network throughput calculated as the sum of the average data rate of all the users.
- J_{Net} : Jain's fairness index for user throughput and is computed as $(\sum_{i=1}^N R_i)^2 / N \sum_{i=1}^N R_i^2$, where R_i is the average user throughput during T . Fairness is computed based on the total data a user receives during T since we are interested long-term fairness.
- VD_{Net} : the average percentage of network VD of all users during T , and is computed as $\text{VD}_{\text{Net}} = \frac{1}{NT} \sum_{\Delta t=1}^T \sum_{i=1}^N \text{VD}_{i,t+\Delta t} / \tilde{D}_{i,t+\Delta t}$.

C. Full Buffer Traffic

Fig. 5 show the results of PRA-MaxNetRate and PRA-MaxMin for full buffer user traffic. In Fig. 5(a) we see that T_{Net} of PRA-MaxNetRate converges to the reference MR allocator, which is expected, as with full buffer traffic the throughput cannot be increased further. Fig. 5(b) and Fig. 5(c) compare the fairness levels of the allocators in the road network and RWP mobility respectively. The figures show how PRA-MaxMin ensures a Jain's fairness index of 1 for both mobility scenarios which is a result of the optimization objective in (2). The fairness gain over other allocators is larger in the RWP scenario as users follow routes that have large variances in the average SNR values, resulting in poor fairness. Also note in Fig. 5(a) that PRA-MaxMin has a high throughput in addition to the high fairness.

D. File Download Traffic

The throughput gains of PRA-MaxNetRate with FTP traffic are now apparent in Fig. 6(a), as opposed to the full buffer case where no increase in T_{Net} was possible. As the number of users increases, we see an increase in T_{Net} compared to the MR allocation scheme. This is because PRA-MaxNetRate can delay serving a user in a congested cell, if the user is moving to a lower density, high data-rate region of the network. When users have finite data requests this long-term RA planning can increase overall network throughput. The figures also show that the throughput of PRA-MaxMin can be higher than MaxRate, while simultaneously achieving a very high long-term fairness as shown in Fig. 6(b) and Fig. 6(c). As the file download size increases, the performance of the PRA schemes will approach the case of full buffer traffic seen in Fig. 5.

E. Buffered Video Streaming

Fig. 7(a) and Fig. 7(b) illustrate the simultaneous improvements in T_{Net} and VD_{Net} achieved by the optimal pre-allocations made in PRA-VDmin. The extent of VD_{Net} gains is even larger with RWP mobility as observed in Fig. 7(c), where at 50 users, VD is 1% for PRA-VDmin, while it is greater than 6% for the remaining allocators. We also see that PRA-VDmin can support close to 100 users at a VD of 6%, thereby doubling the number of supportable users at this VD







Fig. 5. Network throughput T_{Net} and Jain's Fairness Index J_{Net} vs. number of users N for full buffer traffic.




(a) Network throughput for RWP mobility




(b) Jain's fairness index for RWP mobility




(c) Jain's fairness index for the road network

Fig. 6. Network throughput T_{Net} and Jain's Fairness Index J_{Net} vs. number of users N for file download traffic.

(a) Network throughput for the road network



(b) Video Degradation for the road network



(c) Video Degradation for RWP mobility

Fig. 7. Network throughput T_{Net} and Video Degradation VD_{Net} vs. number of users N for buffered video streaming.

level. It is worth pointing that Max-Rate outperforms Rate-Proportional at high load in Fig. 7(b), as in the road network, users have bell shaped SNR variations with time, and therefore Max-Rate does not lead to user starvation. This is not the case with RWP mobility, where Max-Rate performs very poorly.

VI. CONCLUSION

In this paper, we presented a Predictive Resource Allocation (PRA) framework that exploits user rate predictions to improve network and user QoS. We formulated PRA over a network of multiple cells for three different objectives as linear programs. Numerical results of the optimal solutions were presented to provide performance benchmarks in a practical road network, and for the more general RWP mobility scenario. Significant QoS gains over reference RA schemes that do not utilize any predictions were observed.

Although PRA was formulated as LPs, which have reasonable complexity, the dimension of these problems is large due to the long-term planning horizon and multiple cells involved. As this results in significant computation time and memory requirements, future work includes developing real-time algorithms that achieve close to optimal performance, and distributed schemes operating independently at BSs for minimal signalling overhead. Studying the effect of prediction errors and the development of algorithms robust to such errors is another needed area of research.

REFERENCES

- [1] OpenSignal, "Opensignal, project homepage." <http://opensignal.com/>, 2013. Link verified on Mar. 15th, 2013.
- [2] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, "Understanding individual human mobility patterns," *Nature*, vol. 453, pp. 779–782, 2008.
- [3] G. Song, Y. Li, G. Song, and Y. Li, "Utility-based resource allocation and scheduling in ofdm-based wireless broadband networks," *IEEE Communications Magazine*, vol. 43, pp. 127 – 134, Dec. 2005.
- [4] M. Katoozian, K. Navaie, and H. Yanikomeroglu, "Utility-based adaptive radio resource allocation in OFDM wireless networks with traffic prioritization," *IEEE Trans. Wireless Commun.*, vol. 8, no. 1, pp. 66–71, 2009.
- [5] E. Hossain and Z. Lei, eds., *Multicell Cooperation*, *IEEE Wireless Commun.* Feb. 2013.
- [6] H. Abou-zeid, H. Hassanein, and S. Valentin, "Lookback scheduling for long-term quality-of-service over multiple cells," in *Proc. QoS and QoE in Wireless Commun./Networks Workshop, collocated with IEEE IWCMC, to appear*, July 2013.
- [7] S. H. Ali, V. Krishnamurthy, and V. C. M. Leung, "Optimal and approximate mobility-assisted opportunistic scheduling in cellular networks," *IEEE Trans. Mobile Comput.*, vol. 6, pp. 633–648, June 2007.
- [8] J. Yao, S. Kanhere, and M. Hassan, "Improving QoS in high-speed mobility using bandwidth maps," *IEEE Trans. Mobile Comput.*, vol. 11, pp. 603–617, Apr. 2012.
- [9] H. Abou-zeid, S. Valentin, and H. Hassanein, "Context-aware resource allocation for media streaming: Exploiting mobility and application-layer predictions," in *Proc. Capacity Sharing Workshop*, Oct. 2011.
- [10] M. Behrisch, L. Bieker, J. Erdmann, and D. Krajzewicz, "SUMO - Simulation of Urban MObility: An Overview," in *Proc. Third Int. Conf. on Advances in System Simulation (SIMUL 2011)*, pp. 63–68, Oct. 2011.
- [11] 3GPP, "LTE; E-UTRA; radio frequency system scenarios," Technical Report TR 36.942 V11.0.0, 3GPP, 2012.