Cellular Multi-Coverage with Non-uniform Rates

Omer Gurewitz, Yakov Sandomirsky and Gabriel Scalosub
Department of Communication Systems Engineering
Ben-Gurion University of the Negev
Beer-Sheva 84105, Israel
Email: {gurewitz,sandomir,sgabriel}@bgu.ac.il

***Abstract*—Recent advances in the standardization of 4G cellular networks introduce the notion of multi-coverage, where multiple base stations may collaboratively satisfy the demands of mobile users. We provide a theoretical model for studying such multi-coverage environments, in highly heterogeneous settings, where users demands and profits may vary, as can base stations' capacities and the rates with which they can service the users.**

Whereas previous works provided solutions that were only applicable to scenarios where rates are uniform throughout the network, or allowed a mobile user to be serviced by at most one base station, we present several algorithms for the multi-coverage problem in the presence of non-uniform rates, and analyze their performance. We complete our study by a simulation study that further validates our results and provides further insight into algorithm design, depending on the users' characteristics.

I. Introduction

Next generation cellular systems, 3GPP *Long Term Evolution* (*LTE*), [1], and IEEE 802.16m (WiMAX), [2] provide integrated broadband services while maintaining diverse *Quality of Service* (*QoS*) guarantees. In these systems, each mobile user is associated with a base station (BS)[1]. The base station allocates (schedules) distinct frequency-time chunks to the active mobile users within its cell. In order to cope with the high traffic demands, the density of base stations is expected to grow exponentially in the near future [3], where each cell is expected to have a different size determined by the maximum range at which the users can successfully hear the transmitter and by the maximum combined data rate of all the mobiles in the cell. Accordingly, coverage will be provided by a mixture of cells, comprising *Macrocells* which provide wide-area coverage of a few kilometers, *Microcells* which cover a few hundred meters and are suitable for densely populated urban areas, combined with smaller cells such as *Picocells* which are installed in large indoor environments such as offices or shopping centers and *Femtocells* which are self-deployed residential base stations, which bring the network closer to users and provide a big leap in performance. The challenge in such dense networks is hence to provide wide coverage supporting high data rate applications anywhere and anytime, and in particular increasing cell edge users' throughput allowing them to also benefit from the high data rate application support.

[1]In this paper we use the terminology of users and base stations which is similar to user equipment (UE) and evolved Node B (eNB), and mobile stations (MS) and base station (BS), in the *LTE* and WiMAX jargon, respectively.

Traditionally each user is associated with a single BS. Accordingly, the BS-user association problem is crucial in multicell networks. Typically a signal to noise ratio approach is used in which a user is associated with the BS whose signal is received with the highest average strength. Sometimes the cell load is also considered, prioritizing cells with lighter traffic loads. Nonetheless, when the traffic loads are high, the system cannot support all users and admission control is required. One should note that in such heavy loads scenarios, when the loads among cells are not balanced, some users may be blocked even though some of their neighboring cells may be not fully utilized. In this paper we evaluate a mechanism which allows more than one base station to transmit data to a single user simultaneously. Note that this suggested mechanism is in the spirit of the current standards which support the coordinated multi-point (CoMP) mechanism between base stations, which not only allows coordination between the BSs but also leverages this coordination and allows more than one base station to transmit data to a single user simultaneously.

Specifically, we consider a multiple base station (BS) network serving a large number of users. Each user has a different PHY data rate from each BS which is a function of the SNR received from this BS. We further assume that different users run different applications with different traffic demands, and are therefore associated with different utilizations and gains. We assume that each user has a minimum traffic demand which if not fulfilled, the user is deemed uncovered. Furthermore, no extra gain can be obtained from allocating extra resources beyond the minimum required by a user, i.e., each user utilization function can be modeled as a step function where the step is located at its minimum demand. We study the user selection problem in highly loaded heterogeneous scheduled access networks.

Typically, the user selection problem considers two closely related problems: (i) admission control, and (ii) schedule and frame building procedure. The admission controller is responsible for selecting the users that can be served. The schedule and frame builder is responsible for the logical positioning of the allocation within the frame. In this study we deal only with the admission control mechanism, which aims at maximizing the overall system benefit under the constraint that the aggregate allocated resources by each BS to its users is feasible. In this paper we overlook the micro scheduling problem and its feasibility, i.e., we assume that any resource distribution by the network which is feasible by each BS individually can be attained (scheduled) system wise. In contrast to common user association problems in which a user receives service from a single BS, in our approach a user can be served by multiple BSs. In particular, in this paper we analyze the achievable gain from allowing a user to be served by multiple BSs. We prove that the multi-association user-selection problem is a hard problem and prove some properties of the multi-association paradigm. In addition, we devise algorithms for solving the problem using this paradigm. We analyze the performance of our algorithms, and provide a simulation study that further highlights the benefits of our approach.

A. Model and Definitions

Consider a bipartite graph $G = (I, J, E)$ where $I = \{1, 2, \ldots, m\}$ is the set of base stations, $J = \{1, 2, \ldots, n\}$ is the set of mobile users, and $E \subseteq I \times J$. Every user $j \in J$ has a non-negative *demand* $D(j)$ (in bits), and a non-negative *profit* $P(j)$, and every base station $i \in I$ has a non-negative *capacity* $T(i)$ (in seconds). Furthermore, every user $j \in J$ and base station $i \in I$ are characterized by a *rate* $R(i, j)$ (in b/s) capturing the service rate station i can provide user j. For every base station $i \in I$ we let $N(i)$ denote the set of users for which $R(i, j) > 0$, and similarly for every user $j \in J$ we let $N(j)$ denote the set of base stations for which $R(i, j) > 0$. We assume without loss of generality that $R(i, j) \leq 1$ for all $i \in I, j \in J$.[2] For every $i \in I, j \in J$ we let $E(i, j) = \frac{P(j)}{D(j)} R(i, j)$ denote the (i, j)-*effectiveness*.

We assume the following: (i) *no-mice*: there exists some $\delta > 0$ such that if $R(i, j) > 0$ then $R(i, j) > \delta$. This assumption essentially means that the minimum rate a user may receive from a base station cannot be infinitesimally small. (ii) *no-hogs*: there exists some $r \in [0, 1]$ such that if $R(i, j) > 0$ then $D(j) \leq rT(i)R(i, j)$. This assumption essentially means that no user will need to use more than an r-fraction of the capacity of a station that can provide him with service.

We further extend our notation to sets of nodes in G such that for every function F defined over all the nodes and some subset $A \subseteq V$, $F(A) = \sum_{v \in A} F(v)$ (if F is real-valued) or $F(A) = \bigcup_{v \in A} F(v)$ (if F is set-valued).

Given a subset of users $S \subseteq J$, a *cover* for S is a function $x : I \times J \to \mathbb{R}^+$ such that the following inequalities hold: (i) for all $i \in I$, $\sum_{j \in S} x_{i,j} \leq T(i)$, and (ii) for all $j \in S$, $\sum_{i \in I} x_{i,j} R(i, j) \geq D(j)$. For any $i \in I$ and $j \in J$, we refer to $x_{i,j}$ as the *service* given by base station i to user j, or simply as the service along the edge (i, j). If $S \subseteq J$ has a cover x, then we say S is *feasible*. Given a cover x for $S \subseteq J$, for every $i \in I$, we let $x(i) = \sum_{j \in J} x_{i,j}$. A cover x for S is said to be a *minimal cover* for S if for any $\varepsilon > 0$ and any $i \in I$ and $j \in J$, x' resulting from reducing $x_{i,j}$ by ε and leaving all other values of x untouched, is no longer a cover for S.

The (r, δ)-*multi-rate cover* ((r, δ)-MRC) *problem* (or simply, MRC) is to find a subset $S \subseteq J$ and a cover x for S such that $P(S)$ is maximized.

[2] This assumption merely serves as a normalizing factor.

B. Our Contribution

In this paper we study the benefits of multi-coverage in 4G cellular networks, in the presence of non-uniform rates. Such capabilities are suggested by recent standardization of 4G technologies such as LTE-Advanced and IEEE 802.16m.

We study the potential benefits of performing multi-coverage compared to the common single-cover approach employed in current cellular technologies. Specifically, we show that even if rates are non-uniform, the difference between the number of users covered by a multi-cover solution and the number covered by a single-cover solution, is at most $m - 1$. Furthermore, we show that the overall profit of a multi-cover is never more than twice the optimal profit of a single-cover solution.

We further devise approximation algorithms and heuristics that aim at maximizing the overall profit gained from fully serviced users. We present a $(1 - r)\delta$-approximation algorithm for the problem, whose approximation ratio is the best known for the values of r and δ satisfying $(1 - r)\delta > 1 - 1/e$ (where e is the base of the natural logarithm). We further devise a heuristic that is inspired by our approximation algorithm.

It is important to note that our analysis and guarantees apply to *arbitrary* profits.

We perform a simulation study examining the various aspects addressed by our work, where we compare the performance of classical single-cover solutions, our approximation algorithm, and our proposed heuristic. Our results show that the multi-cover paradigm yields a significant boost in attained profit.

C. Related Work

User association policies and cell load balancing were broadly studied under different models and assumptions. For example some studies formulated the load balancing problem as a convex optimization problem under various utility functions, e.g., jointly optimized partial frequency reuse and load-balancing [4], fairness criteria [5], [6], and sum rate of all user maximization [7]. A game theoretic approach was taken in [8]. Call admission control in 3GPP networks was also widely studied, e.g., [9]–[11]. Significant work has been done on scheduling algorithm in OFDMA system, e.g., [12], [13].

Our work is closely related to assignment problems studied within the operations research community, and in particular variants of the General Assignment Problem (GAP) which generalize Multiple Knapsack Problems. In the MAX-GAP problem we are given a set of jobs J and a set of machines I such that each machine is associated with some size capacity, and each (job,machine) pair is associated with some profit for assigning the specific job to the specific machine, and some size which would be used up by the machine if the job is assigned to the machine. The goal is to find an assignment of some subset of jobs to machines that maximizes the overall profit while not violating the constraint imposed by each machines' size capacity. A 2-approximation algorithm for the minimization version of the problem was presented in [14], which was later used to devise a 2-approximation for MAX-GAP [15]. This latter work also showed that the GAP problem is APX-hard (hardness later improved in [16]). Generalizations (and algorithms for these generalizations) include Separable Assignment Problems (SAP) studied in [17], where they also provide a $(1-1/e)$-approximation algorithm for MAX-GAP. The current best approximation ratio for GAP is due to [18] which guarantees to generate a $(1-1/e+\varepsilon)$-approximate solution for some small constant ε. Essentially all these algorithms are based on solving linear programs, where for the most part they are based on meticulously rounding the fractional solution to the linear programs. A combinatorial $(2+\varepsilon)$-approximation (for any $\varepsilon > 0$) was provided in [19]. The main difference between our work and this body of research (using GAP terminology) is that we study the benefits of, as well as algorithms for, exploiting the ability to satisfy the packing requirement of every job by splitting this requirement among several machines. Several works [20], [21] considered splittable-assignment of objects to bins, with added overhead (depending on the specific model). Our model allows for such splittable-assignment but incurs no overheads, thus making it generally applicable.

The work most closely related to ours is [22] which studies multi-coverage solutions with uniform rates. They show a $(1-r)$-approximation algorithm for the problem, where their algorithm heavily relies on maximum-flow subproblems. Our problem is considerably more complex, since having non-uniform rates eliminates the ability to rely on maximum-flow solutions. Furthermore, the solutions proposed for uniform rates are inapplicable to the non-uniform rates settings, since it is impossible to naively "trade" coverage. This difficulty lays at the core of our results in Section II-A.

Paper Organization: In Section II we study the relationship between multi-cover solutions, and single-cover solutions. We focus on both structural aspects of such solutions, as well as on the potential profit that can be accrued from each approach. In Section III we consider the approximability of the multi-cover problem in environments with non-uniform rates. We present our approximation algorithm and analyze its performance. We also propose a refined heuristic based on our analytic viewpoint. Finally, in Section IV we present our simulation results, where we compare the performance of several solutions to the problem (both multi-cover and single-cover) and extract further insight as to their benefits and limitations.

II. MULTI-COVERAGE VS. SINGLE-COVERAGE

A. The Structure of a Multi-coverage Solution

In this section we study some fundamental properties of a multi-coverage solution. In particular, we show the following theorem:

Theorem 2.1: Any minimal cover x for a set $S \subseteq J$ can be transformed into a minimal cover x' for S such that at most $(m-1)$ users in S are covered by more than one base station.

We note that in the case of *uniform rates* the above theorem easily follows by applying simple cycle-cancelation arguments on the bipartite graph. However, in the case of non-uniform rates, such simple arguments do not suffice (this was already mentioned in the Section I, and will further be discussed in Section III). The reason why these simple arguments fail is due to the fact that although the capacity of a BS (in seconds) might be swapped between users, this does not translate to equal quantities of service (in bits) that are swapped between users, due to the heterogeneous rates.

The proof of the theorem is based on the following lemma which shows we can cancel cycles of allocations in a cover, even when rates are non-uniform.

Lemma 2.2: If x is a cover for $S \subseteq J$ and there exists a cycle $C = (e_1, \ldots, e_{2k})$ in E such that $x(e_\ell) > 0$ for all $\ell = 1, \ldots, 2k$, then there exists a cover x' for S that agrees with x on all edges in $E \setminus C$, and for which at least one edge $e \in C$ has $x'(e) = 0$.

Proof: Since we consider edges in $E \subseteq I \times J$, we will refer to any edge $e = (i, j)$ as connecting base station i to user j. Assume without loss of generality that for every $\ell = 1, \ldots, k$ we have $e_{2\ell-1} = (\ell, \ell)$, for every $\ell = 1, \ldots, (k-1)$ we have $e_{2\ell} = (\ell+1, \ell)$, and finally $e_{2k} = (1, k)$. See Figure 1 for the schematics of the above cycle. In what follows we will alternately increase and decrease the the amount of resources allocated by x along the cycle C, while trying to maintain that the resulting allocation x' is indeed a cover for S. Assume we wish to decrease $x_{1,1}$ by an amount τ. It can be shown by induction on $i = 1, \ldots, k$ that the following must hold if the resulting allocation x' is to be a cover for S

1) for every $\ell = 1, \ldots, k$, $x_{\ell,\ell}$ must decrease by $\tau \cdot \left[\prod_{h=1}^{\ell-1} \frac{R(h,h)}{R(h+1,h)}\right]$,[3]
2) for every $\ell = 1, \ldots, (k-1)$, $x_{\ell+1,\ell}$ must increase by $\tau \cdot \left[\prod_{h=1}^{\ell} \frac{R(h,h)}{R(h+1,h)}\right]$, and
3) $x_{1,k}$ must increase by $\tau \cdot \left[\prod_{h=1}^{k-1} \frac{R(h,h)}{R(h+1,h)}\right] \cdot \frac{R(k,k)}{R(1,k)}$.

The above amounts are essentially those compensating each user for the increase or decrease in service it receives from its adjacent base stations along the cycle. We refer the reader to Figure 1.

It follows that for every base station save perhaps the base station 1, the capacity constraint is satisfied since the any service increase is compensated by an equal service decrease. It remains to consider base station 1. Note that the increase in service required from base station 1 is $\tau \left[\prod_{h=1}^{k-1} \frac{R(h,h)}{R(h+1,h)}\right] \frac{R(k,k)}{R(1,k)}$, whereas the decrease in its service is τ. In case $\left[\prod_{h=1}^{k-1} \frac{R(h,h)}{R(h+1,h)}\right] \frac{R(k,k)}{R(1,k)} > 1$ such an assignment of service might not be feasible since the capacity constraint of base station 1 might be violated. However, if we reverse the direction of our construction and start by an *increase* of τ to $x_{1,1}$, we obtain the same expressions as above, except for the inverted roles of the edges: an edge along which we previously increased the value of x would now have its value decrease, and vice-a-versa. In such a case we would obtain that the overall decrease is $\tau \left[\prod_{h=1}^{k-1} \frac{R(h,h)}{R(h+1,h)}\right] \frac{R(k,k)}{R(1,k)}$ which is greater than the increase of τ, in which case no violation

[3] We take the empty product to equal 1.

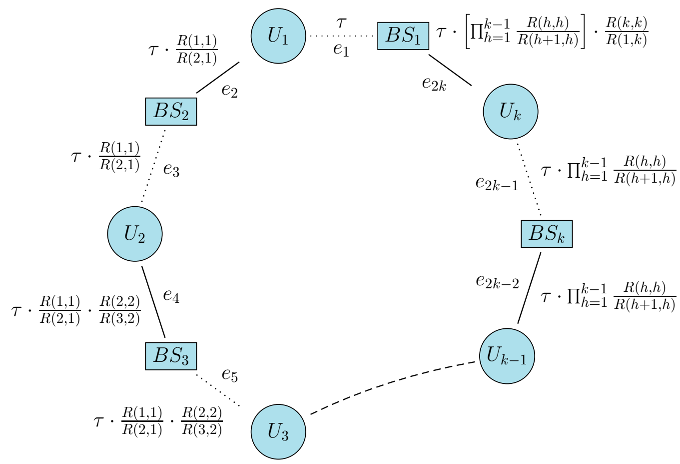

Fig. 1. The schematic structure of a cycle defined by a cover x for $S \subseteq J$. Service is decreased along dotted edges, and increased along solid edges. Each edge is marked with the amount of service increased/decreased along the edge.

of the capacity occurs. It follows that we are guaranteed to have a *direction* along the cycle (either clockwise or counter-clockwise) for which there will be no violation of capacity in any of the stations along the cycle. Assume without loss of generality that the original construction (counter-clockwise in Figure 1) does not violate the capacity of the base stations. Consider all edges along which we are required to decrease the service. By our assumption, in this case those are the edges of type (ℓ, ℓ), $\ell = 1, \dots, k$. In particular, note that the decrease along any such edge (ℓ, ℓ) cannot be greater than the initial amount of service $x_{\ell,\ell}$ provided on that edge, which yields the following set of inequalities:

$$\tau \left[\prod_{h=1}^{\ell-1} \frac{R(h,h)}{R(h+1,h)}\right] \leq x_{\ell,\ell} \quad \ell = 1, \dots, k.$$

This implies a set of k upper bounds on the value of τ (essentially the quotient of dividing $x_{\ell,\ell}$ by its respective product, since these are all fixed given x). Note that taking $\tau = \tau_{\min}$ to be the minimum of these upper bounds implies that at all the inequalities are satisfied, and at least one of them holds with equality. Hence, reducing $x_{1,1}$ by $\tau_{\min}$ would cause the respective edge $e_{\min}$ to have $x'(e_{\min}) = 0$, as required. Note that by our construction every user in S continues to have its demand satisfied, albeit using a different cover. ■

Applying Lemma 2.2 repeatedly, we obtain the following corollary:

Corollary 2.3: Any minimal cover x for a set $S \subseteq J$ can be transformed into a minimal cover x' for S such the set of edges (i, j) for which $x'_{i,j} > 0$ forms a tree.

We can now turn to prove Theorem 2.1.

Proof of Theorem 2.1: Assume x is a minimal cover for $S \subseteq J$. By Corollary 2.3 there exists a minimal cover x' for S such that the set of edges $(i, j) \in I \times J$ for which $x'_{i,j} > 0$ forms a tree. Such a tree spans at most $m + |S|$ vertices and therefore has at most $m + |S| - 1$ edges. Since each user $j \in S$ can be associated with a unique edge connecting it to some base station, it follows that there remain at most $(m-1)$ additional edges in the tree that contribute to the service of multi-covered users. Since each such edge is connected to a single user, the result follows. ■

B. The Profit of a Multi-coverage Solution

When studying the potential profit of using a multi-coverage solution, one of the first question one is faced with is that of the added value such a solution has compared to traditional single coverage solutions. By considering the work of [14], [15], and more specifically the standard integer programming formulations of MAX-GAP, it is straightforward to show that the linear programming relaxations of these formulations of MAX-GAP and that of MRC are the same (more details on these relaxations can be found in Section III).

The analysis of [14], [15] shows that their single-coverage solution is within a factor of 2 from the *optimal fractional solution* of the relaxation. This implicitly implies that the value of an optimal multi-cover solution is never better than twice the value of an optimal single-cover solution. It should be noted that simple toy examples where this gap is almost tight can be easily constructed.

III. APPROXIMATING THE MRC PROBLEM

A. Hardness

Since the MRC problem with non-uniform rates is a generalization of the case where rates are uniform, the hardness results of [22] apply. In particular, if we allow users' demands to exceed the capacity of base stations that can provide them with service (i.e., we consider instances that are not required to satisfy the no-hogs assumption for any $r \leq 1$), then the problem is as hard to approximate as the Maximum Independent Set (MIS) problem, i.e., for any $\varepsilon > 0$ it cannot be approximated to within a factor of $|J|^{\varepsilon}$, unless $\mathrm{NP} = \mathrm{ZPP}$. Furthermore, the problem remains NP-hard for any $r \leq 1$.

B. The Local-Ratio Technique

The algorithm we propose in the sequel is based on the *local-ratio* technique, and is a generalization of the framework set in [22] for uniform rates. It should be noted the dealing with non-uniform rates renders the problem significantly more complex, since computing maximum flow, along with mechanisms of flow augmentation, are no longer applicable when dealing with non-uniform rates. Furthermore, when applying the local-ratio technique one has to take caution when defining the sub-instances in every iteration. We discuss these aspects in further detail in the appropriate places throughout the proof.

As mentioned earlier, the local-ratio technique is based on a decomposition of the profit function into a linear combination of two profit functions. This decomposition is done recursively, where in each level of the recursion some "maximal" solution is returned.[4] Using an inductive argument one shows that the maximal solution approximates well both functions, and hence also their linear combination. More formally, we use

[4]The exact meaning of "maximality" depends on the specific problem, and will be defined shortly.

the following formulation of the local-ratio lemma, adapted to our settings:

Lemma 3.1 (Local-Ratio): Given an instance I to the (r,δ)-MRC problem with profit function p, consider a linear decomposition of p into two objective functions p_1 and p_2 s.t. $p_1+p_2=p$. If x is a cover for some $S\subseteq J$ such that S is an α-approximate solution w.r.t. p_1 and also an α-approximate solution w.r.t. p_2, then S is an α-approximate solution w.r.t. p.

The proof follows from the linearity of the objective function, and its full derivation can be found in, e.g., [23].

In our proof, we decompose the profit function into two components. It is important to note that although one of the components is proportional to the demand, there is no restriction on the original profit function, and in particular, it need not be related to the demand at all.

C. Notions of Maximality

We now turn to define the notion of maximality of a cover x for a set $S\subseteq J$. Given any $A\subseteq I$ and $B\subseteq J$, we define the mixed integer linear program $MILP_{A,B}$, whose relaxation and restrictions will serve as basic building blocks in our algorithm.

$$(MILP_{A,B})\quad \begin{aligned} \max\ & \sum_{j\in B} P(j)y_j & & \\ \text{s.t.}\ & \sum_{j\in B} x_{i,j}\le T(i) & \forall i\in A\quad & (1)\\ & \sum_{i\in A} R(i,j)x_{i,j}\ge D(j)y_j & \forall j\in B\quad & (2)\\ & y_j\in\{0,1\} & \forall j\in B\quad & (3)\\ & x_{i,j}\ge 0 & \forall i\in A, j\in B\quad & (4)\end{aligned}$$

Note that $MILP_{I,J}$ is the mixed integer linear program formulation of the (r,δ)-MRC problem.

By replacing constraint (3) with the constraints

$$\begin{aligned} y_j&\le 1 & \forall j\in B\quad & (5)\\ y_j&\ge 0 & \forall j\in B\quad & (6)\end{aligned}$$

we obtain a linear program, $LP_{A,B}$.

It should be noted that the feasibility of any $S\subseteq J$ can be easily tested by using $LP_{I,J}$ where we require all the constraints (2) and (5) corresponding to $j\in S$ to hold with equality. We further extend the notion of feasibility and say $S\subseteq J$ is feasible for the linear program if this requirement holds.

Given any $S\subseteq J$, we let $\overline{S}=J\setminus S$, and consider the sets $N(\overline{S})$ (the set of base stations that can provide service to some user in $\overline{S}$) and $Y_S=I\setminus N(\overline{S})$ (the set of base stations that have all of their neighbors strictly in S).

For any $S\subseteq J$, we say a a cover plan x for S *drains* Y_S if there exists an *optimal* solution $\tilde{x},\tilde{y}$ for $LP_{Y_S,S}$ such that $x(i)\ge\tilde{x}(i)$ for all $i\in Y_S$.

We can now define our notion of maximality. A subset S with a cover plan x is said to be *LP-maximal* if the following conditions hold: (i) *set-maximality*: for every $j\in\overline{S}$ there is no cover plan for $S\cup\{j\}$, and (ii) *drain-maximality*: x drains Y_S.

As specified in the informal description of our algorithm, we will constantly extend partial solutions to maximal ones. To this end, we need to demonstrate how such a maximal solution can be obtained, and further show that the notion of maximality is sound, i.e., that it does not restrict the feasible solutions of the MRC problem.

Consider the following linear program $\overline{LP}_{A,B}$.

$$(\overline{LP}_{A,B})\quad \begin{aligned} \max\ & \sum_{i\in A, j\in B} E(i,j)x_{i,j} & & \\ \text{s.t.}\ & \sum_{j\in B} x_{i,j}\le T(i) & \forall i\in A\quad & (7)\\ & \sum_{i\in A} R(i,j)x_{i,j}\le D(j) & \forall j\in B\quad & (8)\\ & x_{i,j}\ge 0 & \forall i\in A, j\in B\quad & (9)\end{aligned}$$

Lemma 3.2: $LP_{A,B}$ is equivalent to $\overline{LP}_{A,B}$.

Proof sketch: The equivalence follows from the assignment $y_j=\sum_{i\in A}R(i,j)x_{i,j}/D(j)$ for every $j\in B$. ■

We now turn to show that it suffices to consider feasible solutions S which drain Y_S.

Given some optimal solution $\tilde{x},\tilde{y}$ to $LP_{Y_S,S}$, consider $LP_{I,J}$ with the additional two constraints

$$\begin{aligned} \sum_{j\in S} x_{i,j}&\ge\sum_{j\in S}\tilde{x}_{ij} & \forall i\in Y_S\quad & (10)\\ \sum_{i\in I} R(i,j)x_{i,j}&\ge D(j) & \forall j\in S,\quad & (11)\end{aligned}$$

and denote the resulting linear program by $LP_{I,J}^{Y_S,S}$.

Lemma 3.3: For any $S\subseteq J$, S is feasible for $LP_{I,J}$ if and only if S is feasible for $LP_{I,J}^{Y_S,S}$.

Proof: By definition, any S that is feasible for $LP_{I,J}^{Y_S,S}$ with cover x, is also feasible for $LP_{I,J}$ with the same cover x. Assume S is feasible for $LP_{I,J}$, with a cover x. Define

$$x'_{ij}=\begin{cases} x_{i,j}T(i)/x(i) & \text{if } i\in Y_S \text{ and } x(i)>0\\ T(i)/\left|N(i)\right| & \text{if } i\in Y_S \text{ and } x(i)=0\\ x_{i,j} & \text{otherwise}\end{cases}$$

Note that since for each $i\in I, j\in J$, $x'_{ij}\ge x_{i,j}$, all the constraints (2) are still satisfied, and in particular, since x was a cover for S, so is x'. Furthermore, note that for each $i\in Y_S$ such that $x(i)>0$ we have

$$x'(i)=\sum_{j\in J}x'_{ij}=\sum_{j\in J}x_{i,j}T(i)/x(i)=T(i),$$

for each $i\in Y_S$ such that $x(i)=0$ we have $x'(i)=T(i)$, and for all other i-s we have $x'(i)=x(i)$. This implies that all the constraints (1) are also satisfied. It follows that x',y are a feasible solution for $LP_{I,J}$, and x' is a cover for S. Given any optimal solution $\tilde{x},\tilde{y}$ to $LP_{Y_S,S}$, note that for each $i\in Y_S$, $x'(i)=T(i)\ge\tilde{x}(i)$, which by definition implies that x' satisfies constraint (10). It follows that S is feasible for $LP_{I,J}^{Y_S,S}$ with cover x', thus completing the proof. ■

Corollary 3.4: For any $S \subseteq J$, S has a cover plan if and only if S has a cover plan that drains Y_S.

We conclude our discussion of LP-maximality by noting that testing whether a set of users S has a cover that drains Y_S can be done by solving $LP_{Y_S,S}$ (producing a solution $\tilde{x}$) followed by testing $LP_{I,J}^{Y_S,S}$ for a feasible solution with $\tilde{x}(i)$ as a lower bound on the utilization of each base station i. A feasible solution x for the latter linear program exists if and only if S with cover x drains Y_S. By iteratively testing the feasibility of $S \cup \{j\}$ for every $j \notin S$, we obtain an LP-maximal solution.

D. *The* Rec-MRC *Algorithm*

We now turn to formally define our algorithm for solving the MRC problem. Our algorithm is a recursive algorithm, that uses solutions to linear programs throughout its execution. Given a polynomial size linear program LP with variables z, $\mathrm{Solve}(LP)$ returns an optimal solution z^* for LP if LP has a feasible solution, and zero otherwise. For any (partial) cover x (potentially null-valued), $\mathrm{Full}(x)$ returns the set of users for which constraint (2) is tight, i.e., the set of users S such that x is a cover for S. If x is null, $\mathrm{Full}(x)$ returns the empty set. Our algorithm receives as input a tuple (I, J, P, D, T, R), consisting of the set of base stations, the set of users, the user profit function, the user demand function, the base station capacity function, and the rate matrix, respectively.

Algorithm 1 Rec-MRC(I, J, P, D, T, R)

1: **if** $J = \emptyset$ **then** ▷ recursion base
2: return null
3: **end if**
4: **if** there exists a $j \in J$ such that $P(j) = 0$ **then**
5: $x \leftarrow$ Rec-MRC$(I, J \setminus \{j\}, P, D, T, R)$ ▷ remove j
6: return x
7: **else**
8: for every $j \in J$, set $E(j) = \frac{P(j)}{D(j)}$
9: set $j^* = \arg\min_{j \in J} E(j)$
10: for every $j \in J$, set $P_1(j) = E(j^*) \cdot D(j)$
11: set $P_2 = P - P_1$
12: $x \leftarrow$ Rec-MRC(I, J, P_2, D, T, R)
13: **for** every j such that $P_2(j) = 0$ **do**
14: $S \leftarrow \mathrm{Full}(x)$ ▷ x covers S
15: $z \leftarrow \mathrm{Solve}(LP_{I,S\cup\{j\}}^{Y_{S\cup\{j\}},S\cup\{j\}})$
16: **if** $\mathrm{Full}(z) = S \cup \{j\}$ **then** ▷ test feasibility of $S \cup \{j\}$
17: $x \leftarrow z$ ▷ update cover
18: **end if**
19: **end for**
20: return x
21: **end if**

We first establish some observations on the instances generated by the algorithm during the recursive calls performed throughout its execution.

Lemma 3.5: In every recursive call occurring in line 12, the profit function P_2 is non-negative, and there exists at least one user $j \in J$ for which $P_2(j) = 0$.

Proof: Consider the user j^* identified in line 9, and the value $E(j^*)$. By the definition of P_1 in line 10 and the minimality of $E(j^*)$, for every user $j \in J$,

$$P_1(j) = E(j^*)D(j) \leq E(j)D(j) = P(j).$$

which implies that $P_2 = P - P_1$ is non-negative. Furthermore, for $j = j^*$ the above inequality holds with equality, implying that $P_2(j^*) = 0$. ■

We now turn to analyze the performance of Rec-MRC. First, we bound the overall profit obtained by a *fractional* solution of the linear program $LP_{I,J}$, and relate this to the value of an LP-maximal solution. In what follows, given any linear program LP, we let $\mathrm{OPT}(LP)$ denote the value of an optimal solution for LP.

Lemma 3.6: If S is an LP-maximal solution and $P(\cdot) = D(\cdot)$ then $\mathrm{OPT}(LP_{I,J}) \leq \mathrm{OPT}(LP_{Y_S,S}) + T(N(\overline{S}))$

Proof: Let x be an optimal solution for $LP_{I,J}$. Since $P(\cdot) = D(\cdot)$ we have $\underline{E(i,j)} = R(i,j)$. By using the equivalent linear program $\overline{LP}_{I,J}$ we obtain

$$\begin{aligned}
\mathrm{OPT}(LP_{I,J}) &= \sum_{(i,j)\in Y_S\times S} E(i,j)x_{i,j} + \sum_{(i,j)\notin Y_S\times S} E(i,j)x_{i,j} \\
&\leq \mathrm{OPT}(LP_{Y_S,S}) + \sum_{(i,j)\notin Y_S\times S} R(i,j)x_{i,j} \\
&\leq \mathrm{OPT}(LP_{Y_S,S}) + \sum_{(i,j)\in N(\overline{S})\times J} x_{i,j} \\
&\leq \mathrm{OPT}(LP_{Y_S,S}) + \sum_{i\in N(\overline{S})} x(i,J) \\
&\leq \mathrm{OPT}(LP_{Y_S,S}) + T(N(\overline{S})).
\end{aligned}$$

The equality follows from the equivalence of $LP_{I,J}$ and $\overline{LP}_{I,J}$ shown in Lemma 3.2. The first inequality follows from the fact that $\{x_{i,j} | (i,j) \in Y_S \times S\}$ is a feasible solution for $LP_{Y_S,S}$, the second inequality follows from the fact that $R(i,j) \leq 1$ for all i,j, along with the fact that base stations in Y_S cannot contribute to the coverage of users in $\overline{S}$, and the third and fourth inequalities follow from the definitions. The result follows. ■

Lemma 3.7: If S is an LP-maximal solution then $T(N(\overline{S})) < \frac{x(N(\overline{S}),S)}{1-r}$

Proof: Since S is LP-maximal, it holds that for every $j \in \overline{S}$, $S \cup \{j\}$ is not feasible for $LP_{I,J}$. In particular, for every $j \in \overline{S}$ and $i \in N(\overline{S})$ we have $(T(i) - x(i,S))R(i,j) < D(j)$ (since otherwise we could have set $x_{i,j} = T(i) - x(i,S)$ while keeping all other values of x unchanged, and obtain a feasible solution for $S \cup \{j\}$). Since by our assumption that there are no hogs, $D(j) \leq rT(i)R(i,j)$, we obtain

$$(T(i) - x(i,S))R(i,j) < D(j) \leq rT(i)R(i,j).$$

By rearranging and summing over all $i \in N(\overline{S})$, the result follows. ■

Lemma 3.8: If $P(\cdot) = D(\cdot)$ and $S \subseteq J$ with cover plan x is LP-maximal, then $P(S) \geq (1-r)\delta P(\mathrm{OPT})$

Proof: The result follows from the following series of (in)equalities:

$$
\begin{aligned}
P(\mathrm{OPT}) &= D(\mathrm{OPT}) && (12)\\
&\leq \mathrm{OPT}(LP_{I,J}) && (13)\\
&\leq \mathrm{OPT}(LP_{Y_S,S}) + T(N(\overline{S})) && (14)\\
&\leq x(Y_S,S) + T(N(\overline{S})) && (15)\\
&< x(Y_S,S) + \frac{x(N(\overline{S}),S))}{1-r} && (16)\\
&= \frac{1}{1-r}((1-r)x(Y_S,S) + x(N(\overline{S}),S)) &&\\
&< \frac{1}{1-r}(x(Y_S,S) + x(N(\overline{S}),S)) &&\\
&= \frac{1}{1-r}x(I,S) && (17)\\
&\leq \frac{1}{1-r}\cdot\frac{1}{\delta}D(S) && (18)\\
&= \frac{1}{(1-r)\delta}\cdot P(S). && (19)
\end{aligned}
$$

Equalities (12) and (19) follow from having $P(\cdot) = D(\cdot)$. Inequality (13) follows from the fact that $LP_{I,J}$ is a relaxation of the MRC problem, and inequality (14) follows from Lemma 3.6. Inequality (15) follows from the LP-maximality of x and having $P(\cdot) = D(\cdot)$ and the fact all rates are at most 1. Inequality (16) follows from Lemma 3.7. Equality (17) follows from definition, and finally inequality (18) follows from the fact all rates are at least δ. ■

We now turn to prove the upper bound on the performance of algorithm Rec-MRC, which holds for any profit function.

Theorem 3.9: Algorithm Rec-MRC produces a $(1-r)\delta$-approximate solution.

Proof: We prove by induction on the solution returned by the recursive calls that the solution returned is a $(1-r)\delta$-approximation. For the base case, when $J = \emptyset$, both the optimal solution and the solution returned in line 2 are the empty solutions. It follows that the solution returned is optimal, and in particular a $(1-r)\delta$-approximation.

For the inductive step, there are two cases to consider. The first case is the solution returned in line 6. Since $P(j) = 0$, if x is a $(1-r)\delta$-approximate solution w.r.t. $J \setminus \{j\}$, then it is also a $(1-r)\delta$-approximate solution w.r.t. J. The second case to consider is the solution returned in line 20. We will show that this solution is a $(1-r)\delta$-approximate solution w.r.t. both profit functions P_1 and P_2, and therefore, by lemma 3.1 it will also be a $(1-r)\delta$-approximate solution w.r.t. $P = P_1 + P_2$.

We begin by considering P_2. By the induction hypothesis, the solution x obtained in line 12 is a $(1-r)\delta$-approximate solution w.r.t. profit function P_2. Since every user j added when extending the solution in lines 13–18 has $P_2(j) = 0$, the solution returned in line 20 is a $(1-r)\delta$-approximate solution w.r.t. profit function P_2.

We now turn to consider P_1. Since in lines 13–18 we iteratively try to add users to the solution and provide a witness cover that drains Y_S in each iteration, the solution returned in line 20 is both drain-maximal and set-maximal. By Lemma 3.8 it follows that the resulting solution is a $(1-r)\delta$-approximate solution w.r.t. P_1.

As mentioned earlier, since the solution returned in line 20 is $(1-r)\delta$-approximate w.r.t. both P_1 and P_2, by Lemma 3.1 we conclude that it is $(1-r)\delta$-approximate w.r.t. $P = P_1 + P_2$, which concludes the proof. ■

Our algorithm was presented in a recursive fashion since this adheres more readily to our analysis approach. However, the algorithm can be shown to be equivalent to the following simple greedy approach: Sort the users in decreasing order of their profit-to-demand ratio, and try to add them to the solution one-by-one, in this order. One can add the current user j to the current solution S, if there exists a feasible solution to $LP_{I,S\cup\{j\}}$. In terms of time complexity, if we denote by F the time it takes to solve a the linear program of $LP_{I,J}$ in each iteration, then the overall running time of the algorithm is $O(n \log n + nF)$, where F can be made polynomial by using the ellipsoid method or interior point method [24].

Next, we build upon the algorithmic approach set forth in this section, and more specifically upon the iterative formulation of our proposed algorithm, and design a more refined heuristic for performing the task of multi-coverage.

E. A Refined Greedy Heuristic

The algorithm described in Algorithm 1 does not take into account the available rates for each user, and makes greedy decisions based on the profit-to-demand ratio alone. However, the demand of a user from a base stations is strongly mitigated by the *rate* with which the base station can serve a user. This essentially means that for each base station i and for each user j one can identify the rate-effective profit-to-demand ratio of user j compared to station i as the quantity $\frac{P(j)}{D(j)/R(i,j)} = E(i,j)$ where the denominator essentially captures the actual amount of resources user j would require from base station i, if covered by base station i alone.

This gives rise to the following greedy heuristic, described in Algorithm 2, which aims to refine the greedy approach manifested in the algorithms of Section III. In what follows, for every $X \subseteq J$ and rate matrix R' we denote by $LP_{I,X}(R')$ the solution to the linear program $LP_{I,X}$ using the rate matrix R' (which might be different than the original matrix R defined as part of the input to the problem).

IV. Simulation Study

In order to evaluate the two algorithms suggested in III, we apply them on a variety of network topologies and compare them to the single coverage suggested in [15]. We compare all three heuristics to the fractional solution.

A. Simulation Settings

In this section we describe the settings of our simulations, aimed to examine the performance and benefits of multi-cover solutions to the MRC problem. We conducted our simulation in MATLAB, where we consider settings which capture many of the facets of a 4G cellular environment.

Our baseline deployment consists of 8 hexagonal cells. Each cell is associated with a macro-cell base station at its center, and all base stations use the same fixed transmit power of 50W. This serves as our underlying *cellular* network. We consider n users distributed uniformly at random within the region. We

Algorithm 2 Heuristic-MRC(I, J, P, D, T, R)

$S \leftarrow \emptyset$
sort all pairs (i, j) in decreasing order of $E(i, j)$
for every pair $(i, j) \in I \times J \setminus S$ (in order) **do**
  for every $j' \in J$ **do**
    $\text{mr}_{i,j}(j') = \min\{R(i', j') \mid E(i', j') \geq E(i, j)\}$ ▷ set the iteration min-rate for j'
  end for
  $R_{i,j} = \vec{0}$ ▷ initiate iteration rate matrix to zero
  for every $(i', j') \in I \times S \cup \{j\}$ **do**
    if $R(i', j') \geq \text{mr}_{i,j}(j')$ **then**
      $R_{i,j}(i', j') = R(i', j')$ ▷ use only rates above the current min-rate
    end if
  end for
  $x \leftarrow \text{Solve}(LP_{I, S\cup\{j\}}(R_{i,j})$
  if $\text{Full}(x) = S \cup \{j\}$ **then**
    $S \leftarrow S \cup \{j\}$
  end if
end for
return S

use the Shannon capacity for computing the rates between the users and the base stations, where the path losses (in dB) were calculated according to the COST-Hata model [25, Chapter 4, Equation 4.4.3]. For example, the rate at the edge of a cell is 1.7Mbps. Depending on the value of δ, we prune all rates which are below a fraction of δ of the maximum rate supported by a base station, and consider only rates above this threshold. This value affects the underlying bipartite graph describing the connectivity of users to base stations. In the extreme cases, $\delta = 1$ implies a coverage radius of zero for each base station, whereas $\delta = 0$ implies an infinite coverage radius, which in turn results in a complete bipartite graph.

We consider the case of uniform base station time capacities, normalized to 1. Given the parameter r, for any user j and base station i for which $R(i, j) > 0$, we obtain an upper bound on the permissible demand of user j (due to our "no-hogs" assumption). The minimum of these bounds, denoted M_j, serves as an upper bound on the demand of user j. We pick its demand uniformly at random in the range $[M_j/50, M_j]$. The profit of satisfying a user's demand is either chosen uniformly at random from the range $[0, 100]$ (referred to as *random profits*), or they are chosen as equal to the demand, which means the system aims at maximizing the overall throughput.

B. Simulation Results

Figures 3-2 show the results of our simulation study, where we examine the performance of the two algorithms proposed in the previous sections, Rec-MRC and Heuristic-MRC, with the performance of the single coverage algorithm suggested in [15], referred to as CK-GAP. In each of the plots we describe the performance of each algorithm by its *performance ratio*, which is the ratio between the algorithm's performance, and that of the optimal fractional solution to the LP relaxation of the problem. This latter value serves as an upper bound on the optimal performance possible, and therefore serves as a benchmark. Nonetheless, one should bear in mind that since the fractional solution of the LP relaxation is usually infeasible for the integral problem, the actual distance from the optimal solution is usually smaller. For each setting we conducted a set of 100 independent simulations, and the results show the resulting average performance ratio and confidence intervals.

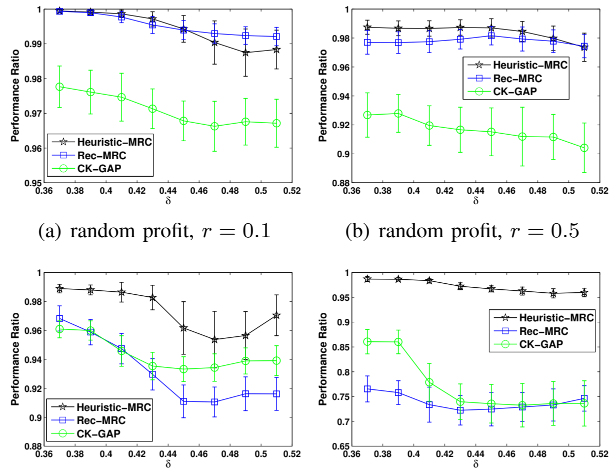

(a) random profit, $r = 0.1$ (b) random profit, $r = 0.5$

(c) profit equals demand, $r = 0.1$ (d) profit equals demand, $r = 0.5$

Fig. 2. Simulation results for varying value of δ

In Figure 2 we explore the effect of a minimum rate restriction on the performance of the various algorithms in a system with $n = 250$ users. In a path-loss-bounded environment the minimum rate restriction is equivalent to limiting the coverage radius of each base station, which is commensurate with the average degree of the users in the underlying bipartite graph. We recall that increasing δ implies a smaller coverage radius, which results in a smaller average degree. In our simulations we let δ vary in the range [0.37,0.51] which is equivalent to allowing a coverage radius between $2R$ and R, where R is the radius of the circumcircle of a single cell. For example, for $\delta = 0.51$ the vast majority of users can only receive service from a single base station. Increasing δ beyond that would result in disjoint knapsack problems where no user can receive service from more than one base station.

In general all algorithms exhibit some performance degradation as we increase δ (and reduce the coverage radius of the base stations). One should note that this degradation is not only absolute, but also relative to the solution of the LP relaxation. For random profits (Figures 2(a)-2(b)), both of our proposed algorithms employing multi-coverage exhibit a performance which is very close to optimal, while the contending single coverage algorithm CK-GAP shows a significantly inferior performance, of up to 10%. Results from additional simulations not presented here show that the gap in performance continues to increase as users demands increase (governed by the parameter r). The situation is more complex in the case where profit equals demand. First, it should be noted that our proposed algorithm Heuristic-MRC outperforms all other algorithms by a significant margin. This can be attributed to

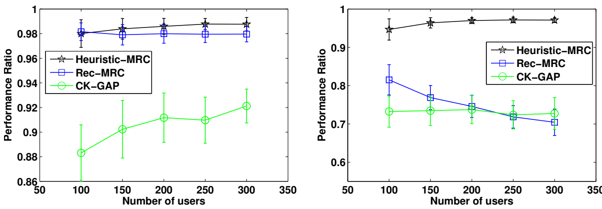

(a) random profit

(b) profit equals demand

Fig. 3. Simulation results for varying number of users

the fact that this algorithms exploits multi-coverage ability, while taking into accounts rate diversity. In the case where $r = 0.1$ (Figure 2(c)), i.e., when demands are relatively small, the benefits of multi-coverage are more limited, and even as δ increases, the performance degradation of CK-GAP is relatively mild. When considering Rec-MRC in this case where profits equal demand, the algorithm essentially adds users in arbitrary order (since they all have a profit-to-demand ratio of 1), and therefore suffers from severe performance degradation as δ increases, and multi-coverage diversity decreases. In the case where $r = 0.5$ (Figure 2(d)), i.e., when demands are more imposing, the single-coverage algorithm CK-GAP is limited in its ability to fully use the capacity of the base stations, leaving a more significant fraction of their capacity unused, and therefore exhibiting fast-deteriorating performance. Algorithm Rec-MRC exhibits poor performance for all values of δ, again – due to its inability to discern between users.

In Figure 3 we consider the effect of the number of users in the system on the performance of the various algorithms, with parameters $r = 0.5$ and $\delta = 0.44$. The results show the performance ratios for both random profits (Figure 3(a)), and the case where profits equal demand (Figure 3(b)). In both cases the performance of Heuristic-MRC and CK-GAP slightly improves as user diversity increases, since both algorithms are able to exploit this diversity. The same applies to Rec-MRC in the case of random profits. However, in the case where profits equal demand, Rec-MRC essentially adds users in arbitrary order (as discussed earlier concerning Figure 2), and therefore shows a steady decrease in performance.

V. Conclusions

We consider the problem of multi-coverage in 4G cellular networks, where rates between base stations and users are non-uniform. We explore the benefits of such an approach compared to commonly used single-coverage solutions, and design algorithms which aim at maximizing the profit obtained from fully satisfied users. We provide upper bounds on the performance of some of our algorithms, and use these analytic insights to design further heuristics for the problem. Finally, we conduct a simulation study which further validates our results, and demonstrates the benefit of multi-coverage.

Acknowledgements

This work was supported by the CORNET consortium, sponsored by the Magnet Program of Israel MOITAL.

References

[1] S. Sesia, M. Baker, and I. Toufik, *LTE-The UMTS Long Term Evolution: From Theory to Practice*. Wiley, 2011.

[2] "IEEE standard for local and metropolitan area networks. Part 16: Air interface for broadband wireless access systems, Amendment 3: Advanced air interface," May 2011.

[3] Informa, "Small cell market status (white paper)," February 2013. [Online]. Available: http://www.informatandm.com/

[4] K. S. S. Chong and G. de Veciana, "Dynamic association for load balancing and interference avoidance in multi-cell networks," *IEEE Transactions on Wireless Communications*, vol. 8, no. 7, pp. 3566–3576, 2009.

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

[6] Q. Ye, B. Rong, Y. Chen, M. Al-Shalash, C. Caramanis, and J. G. Andrews, "User association for load balancing in heterogeneous cellular networks," *IEEE Transactions on Wireless Communications*, vol. 12, no. 6, pp. 2706–2716, 2013.

[7] S. Corroy, L. Falconetti, and R. Mathar, "Dynamic cell association for downlink sum rate maximization in multi-cell heterogeneous networks," in *IEEE ICC*, 2012, pp. 2485–2489.

[8] L. Jiang, S. Parekh, and J. C. Walrand, "Base station association game in multi-cell wireless networks," in *IEEE WCNC*, 2008, pp. 1616–1621.

[9] O. Delgado and B. Jaumard, "Joint admission control and resource allocation with gos and qos in lte uplink," in *IEEE Globecom Workshop on Broadband Wireless Access*, 2010, pp. 829 – 833.

[10] S. Bashar and Z. Ding, "Admission control and resource allocation in a heterogeneous ofdma wireless network," *IEEE Transactions on Wireless Communications*, vol. 8, no. 8, pp. 4200–4210, 2009.

[11] R. Kwan, R. Arnott, and M. Kubota, "On radio admission control for lte systems," in *VTC Fall*, 2010.

[12] R. Cohen and L. Katzir, "Computational analysis and efficient algorithms for micro and macro ofdma downlink scheduling," *IEEE/ACM Transactions on Networking*, vol. 18, no. 1, pp. 15–26, 2010.

[13] Y. Ben-Shimol, I. Kitroser, and Y. Dinitz, "Two-dimensional mapping for wireless ofdma systems," *IEEE Transactions on Broadcasting*, vol. 52, no. 3, pp. 388–396, 2006.

[14] D. Shmoys and E. Tardos, "An approximation algorithm for the generalized assignment problem," *Mathematical Programming*, vol. 62, no. 3, pp. 461–474, 1993.

[15] C. Chekuri and S. Khanna, "A polynomial time approximation scheme for the multiple knapsack problem," *SIAM Journal on Computing*, vol. 35, no. 3, pp. 713–728, 2005.

[16] D. Chakrabarty and G. Goel, "On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP," *SIAM Journal on Computing*, vol. 39, no. 6, pp. 2189–2211, 2010.

[17] L. Fleischer, M. X. Goemans, V. S. Mirrokni, and M. Sviridenko, "Tight approximation algorithms for maximum separable assignment problems," *Mathematics of Operations Research*, vol. 36, no. 3, pp. 416–431, 2011.

[18] U. Feige and J. Vondrák, "Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$," in *IEEE FOCS*, 2006, pp. 667–676.

[19] R. Cohen, L. Katzir, and D. Raz, "An efficient approximation for the generalized assignment problem," *Information Processing Letters*, vol. 100, no. 4, pp. 162–166, 2006.

[20] N. Naaman and R. Rom, "Packet scheduling with fragmentation," in *IEEE INFOCOM*, 2002, pp. 427–436.

[21] H. Shachnai, T. Tamir, and O. Yehezkely, "Approximation schemes for packing with item fragmentation," *Theory of Computing Systems*, vol. 43, no. 1, pp. 81–98, 2008.

[22] D. Amzallag, R. Bar-Yehuda, D. Raz, and G. Scalosub, "Cell selection in 4G cellular networks," *IEEE Transactions on Mobile Computing*, vol. 12, no. 7, pp. 1443–1455, 2013.

[23] R. Bar-Yehuda, "One for the price of two: A unified approach for approximating covering problems," *Algorithmica*, vol. 27, no. 2, pp. 131–144, 2000.

[24] D. Bertsimas and J. N. Tsitsiklis, *Introduction to Linear Optimization*. Athena Scientific, 1997.

[25] E. Damosso, *Digital Mobile Radio Towards Future Generation Systems, COST 231 Final Report*, 1999.