Opportunistic Gains of Mobility in Cellular Data Networks

Nivine Abbas
Orange Labs & Telecom ParisTech
Paris, France
nivine.abbas@telecom-paristech.fr

Thomas Bonald
Telecom ParisTech
Paris, France
thomas.bonald@telecom-paristech.fr

Berna Sayrac
Orange Labs
Issy-Les-Moulineaux, France
berna.sayrac@orange.com

Abstract—**In this paper, we assess the performance gains of mobility on the downlink of cellular data networks. These gains are only due to the elastic nature of traffic and thus observed even under a blind, fair scheduling scheme: data are more likely transmitted when users are close to the base stations, in good radio conditions. This phenomenon is further amplified by opportunistic scheduling schemes that exploit multi-user diversity. The results are based on the analysis of flow-level traffic models and validated by system-level simulations.**

Keywords—*Cellular data networks, mobility, opportunistic scheduling, flow-level modeling, queuing theory.*

I. INTRODUCTION

The efficiency of wireless data transmissions strongly depends on the variations of the radio channel quality, especially the *fast* variations caused by the constructive and destructive combinations of the multiple signal paths between the transmitter and the receiver. While traditional wireless systems use diversity techniques to mitigate these channel variations and thus provide constant bit rates, advanced wireless systems exploit these variations through so-called opportunistic schedulers, taking advantage of the inherent "elasticity" of data transfers to increase the overall system capacity [16], [2], [3].

In this paper, we show that *slow* channel variations, due to the mobility of the device, can be exploited as well. In fact, they are already exploited by current packet schedulers, even blind strategies like the round-robin policy. The reason is that, due to mobility, data are more likely transmitted when users are close to the base stations, in good radio conditions. This phenomenon is further amplified by opportunistic schedulers.

The fact that mobility may improve performance has already been observed in the literature, see for instance [4], [5], [7], [8], [11], [13]. However, these papers mainly give *qualitative* results on the impact of mobility. In the present paper, we assess the intra-cell and inter-cell mobility-induced gain under a fair scheduling scheme (the round-robin policy) and an opportunistic scheduling scheme (the max C/I policy, see [15]) for both the cell center and the cell edge. We show in particular that under both scheduling policies, mobility improves throughput performance at the cell edge. This is a rather surprising result for the max C/I policy, which gives priority to cell-center users. The reason is that these users quickly complete their file transfers, which saves scheduling resources for cell-edge users. The results are based on the analysis of flow-level traffic models and validated by system-level simulations.

The rest of the paper is organized as follows. We present the reference model in the absence of mobility in the next section. We analyse the impacts of intra-cell and inter-cell mobility in Sections III, IV and V. Section VI investigates the impact of fast fading. Section VII presents the system-level simulations and Section VIII concludes the paper.

II. REFERENCE MODEL

We first present the reference model in the absence of mobility and the corresponding performance results under round-robin and max C/I scheduling schemes.

A. Cellular network

We model each cell by a set of N regions. In each region, radio conditions are supposed to be homogeneous and thus users are served at the same physical data rate on the downlink. In the simple case of two regions illustrated by Figure 1, users may be close to the cell center and experience good radio conditions (light gray) or close to the cell edge and suffer from bad radio conditions (dark gray). We model each region by a queue with a specific service rate corresponding to the physical data rate in this region; since all users in the cell share the same radio resources, each cell can be viewed as a set of N parallel queues with coupled processors. The precise coupling depends on the scheduling policy, as explained below.
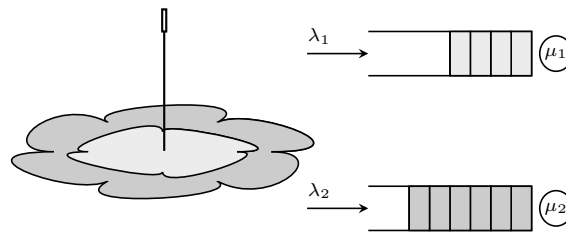

Fig. 1. A simple model with two cell regions.

B. Traffic model

We consider elastic traffic only, corresponding to data transfers. Specifically, we assume that new data flows are generated in region i at the random times of a Poisson process of intensity λ_i. We denote by $\lambda = \sum_i \lambda_i$ the total flow arrival rate in the cell and by $p_i = \lambda_i/\lambda$ the probability that a new data flow is generated by a user in region i.

For the sake of tractability, we use a flow-level model where each data flow is viewed as a fluid of random volume to be transmitted. The volumes have an exponential distribution with mean σ (in bits). When region i is served, flows are completed at rate μ_i in the absence of fast fading, corresponding to the physical rate $\mu_i\sigma$ (in bit/s). We assume regions are numbered in decreasing order of physical rates, that is $\mu_1 > \mu_2 > \ldots > \mu_N$.

Queue i represents the number of active flows in region i. The load of queue i is the ratio $\rho_i = \lambda_i/\mu_i$ of arrival rate to service rate. Since the radio resources are shared in time, the actual service rate of queue i is modulated by ϕ_i, the fraction of time spent by the scheduler on users in region i. This depends on the system state and the scheduling policy. For work-conserving policies, we have $\sum_i \phi_i = 1$. The overall stability condition is then $\rho < 1$ in the absence of mobility, with $\rho = \sum_i \rho_i$. We shall see in Sections III and IV that mobility may increase the stability region. Observe that the cell load may be written $\rho = \lambda/\mu$ where μ is the harmonic mean service rate weighted by the distribution of arrivals in the N regions:

$$\mu = \frac{1}{\sum_i p_i/\mu_i}.$$

C. Throughput metrics

We measure performance in terms of mean throughputs in the different regions. We denote by $X(t)$ the system state at time t. This is an N-dimensional vector whose component i gives the length of queue i at time t. It is an irreducible Markov process with stationary distribution π. In any state x such that $x_i > 0$, each user in region i has throughput $\mu_i\sigma\phi_i(x)/x_i$. Now the distribution seen by users in region i is the size-biased distribution [6]:

$$\pi_i(x) \propto x_i\pi(x).$$

Observe in particular that $\pi_i(x) = 0$ in all states x such that $x_i = 0$. We denote by E_i the corresponding expectation. The mean throughput of users in region i is then given by

$$\gamma_i = E_i\left(\frac{\mu_i\sigma\phi_i(X)}{X_i}\right) = \frac{E(\mu_i\sigma\phi_i(X))}{E(X_i)}. \tag{1}$$

By the traffic conservation equation

$$\lambda_i = E(\mu_i\phi_i(X)), \tag{2}$$

we deduce

$$\gamma_i = \frac{\lambda_i\sigma}{E(X_i)}. \tag{3}$$

Observe that, by Little's law, this is the ratio of mean flow size (in bits) to mean flow duration in region i.

By a similar argument, the mean throughput in the cell is

$$\gamma = \frac{E(\sum_i \mu_i\sigma\phi_i(X))}{E(\sum_i X_i)}. \tag{4}$$

By the traffic conservation equation

$$\lambda = \sum_i E(\mu_i\phi_i(X)), \tag{5}$$

we deduce

$$\gamma = \frac{\lambda\sigma}{\sum_i E(X_i)}. \tag{6}$$

We shall see that, in the absence of mobility, the mean throughput in region i decreases from the physical data rate $\mu_i\sigma$ to 0 when the cell load ρ grows from 0 to 1, while the mean throughput in the cell decreases from the harmonic mean physical data rate $\mu\sigma$ to 0, the precise impact of the cell load depending on the scheduling policy.

D. Round-robin policy

Under the round-robin policy, users share the radio resources equally, independently of their radio conditions. Thus users in region i are allocated a fraction

$$\phi_i(x) = \frac{x_i}{\sum_j x_j}$$

of radio resources in state x. The transition graph of the Markov process $X(t)$ is shown in Figure 4 (omitting the blue arrows, corresponding to mobility). This Markov process is reversible and corresponds to the state of a multi-class processor-sharing queue. The stationary distribution of $X(t)$ is:

$$\pi(x) = (1-\rho)\binom{\sum_i x_i}{x_1,\ldots,x_N}\rho_1^{x_1}\ldots\rho_N^{x_N}.$$

We deduce the mean number of flows in region i:

$$E(X_i) = \frac{\rho_i}{1-\rho},$$

and, in view of (3), the mean throughput in region i:

$$\gamma_i = \mu_i\sigma(1-\rho). \tag{7}$$

Observe that the mean throughputs decrease *linearly* with the cell load in all regions.

In view of (6), the mean throughput in the cell is

$$\gamma = \mu\sigma(1-\rho). \tag{8}$$

It also decreases linearly with the cell load.

E. Max C/I policy

The max C/I policy is an opportunistic scheduling strategy that prioritizes those users with the best radio conditions. For two regions for instance, cell-center users are scheduled first and are allocated all the resources whenever active; cell-edge users are served only when there are no active cell-center users. The transition graph of the Markov process $X(t)$ is shown in Figure 5 (omitting the blue arrows). The Markov process is no longer reversible and corresponds to a preemptive priority queue. Applying known results of queuing theory [14], we get the mean number of flows in region i:

$$E(X_i) = \frac{\rho_i}{1-\sum_{j<i}\rho_j}\left(1+\frac{\rho}{1-\sum_{j<i}\rho_j}\mu_i\mu\sum_{j<i}\frac{p_j}{\mu_j^2}\right)$$

and, in view of (3), the mean throughput in region i:

$$\gamma_i = \mu_i\sigma\frac{1-\sum_{j<i}\rho_j}{1+\frac{\rho}{1-\sum_{j\leq i}\rho_j}\mu_i\mu\sum_{j<i}\frac{p_j}{\mu_j^2}}.$$

Observe that the mean throughput in region i is positive whenever $\sum_{j<i}\rho_j < 1$: it decreases from $\mu_i\sigma$ to 0 when the load generated by users in regions $1,\ldots,i$ grows from 0 to 1.

Similarly, we get the mean throughput in the cell

$$\gamma = \mu\sigma \left(\sum_i \frac{\rho_i/\rho}{1 - \sum_{j<i} \rho_j} \left(1 + \frac{\rho}{1 - \sum_{j \le i} \rho_j} \mu_i \mu \sum_{j<i} \frac{p_j}{\mu_j^2} \right) \right)^{-1}$$

which decreases from $\mu\sigma$ to 0 when the cell load ρ grows from 0 to 1. This is the harmonic mean of the flow throughputs in the N regions.

F. Numerical results

Figure 2 compares the throughput performance of both policies for two regions, with $p_1 = p_2 = 1/2$, $\mu_1 = 8$, $\mu_2 = 2$ and $\sigma = 1$. Observe that, under max C/I scheduling, the cell-center throughput is positive whenever $\rho_1 < 1$, that is $\rho < 5$, and decreases linearly. The throughput gain compared to the round-robin policy is huge for cell-center users, at the expense of a slight throughput degradation for cell-edge users. The cell-average throughput remains approximately the same since it corresponds to the harmonic mean throughput over the cell and thus is mainly affected by cell-edge users.
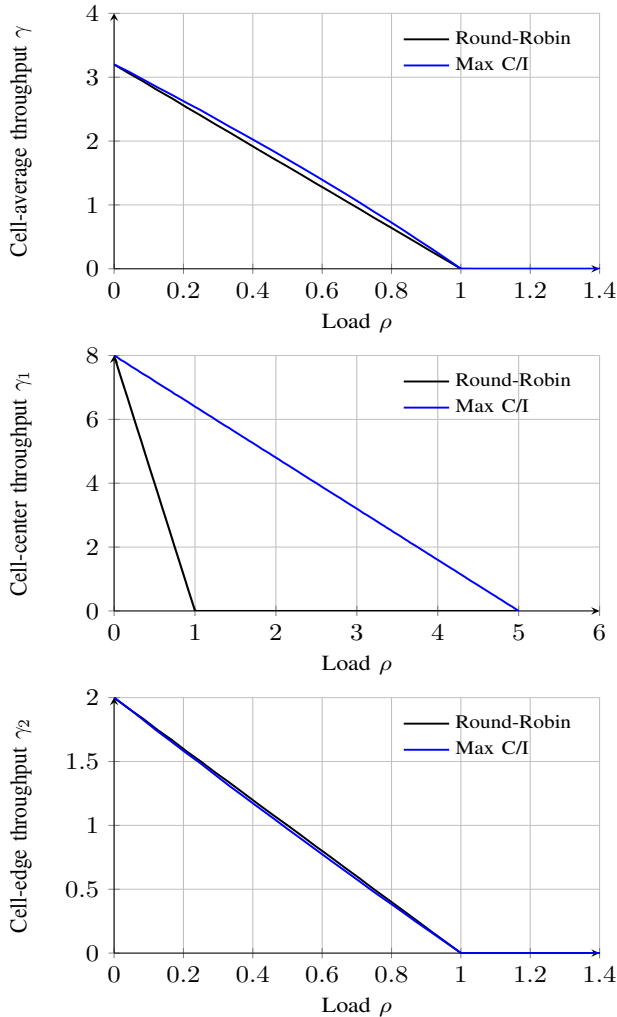

Fig. 2. Throughput performance for two regions without mobility.

III. IMPACT OF INTRA-CELL MOBILITY

In this section, we add intra-cell mobility to the previous model, as shown in Figure 3 for two regions. We suppose that users move from the center of the cell to the edge and vice versa. We still assume that there is no fast fading.
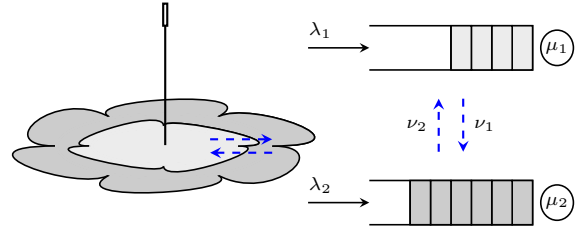

Fig. 3. Queuing model for two regions with intra-cell mobility.

A. Mobility model

We assume that each user in region i moves to region $i-1$ (for $i > 1$) and to region $i+1$ (for $i < N$) after exponential durations, at respective rates $\nu_{i,i+1}$ and $\nu_{i,i-1}$. The probability that a user is in region i then satisfies:

$$q_i \propto \prod_{j=1}^{i-1} \frac{\nu_{j,j+1}}{\nu_{j+1,j}}.$$

Note that this is not the probability that an *active* user is in region i, which is given by

$$p'_i = \frac{E(X_i)}{\sum_j E(X_j)}$$

and depends both on the cell load and on the scheduling policy. In state x, the total mobility rate from region i to region $i+1$ and from region $i+1$ to region i (for $i < N$) are $x_i \nu_{i,i+1}$ and $x_{i+1} \nu_{i+1,i}$, respectively. For notational convenience, we write $\nu_1 \equiv \nu_{1,2}$ and $\nu_2 \equiv \nu_{2,1}$ for two regions. The corresponding state diagrams are given by Figures 4 and 5, where the blue arrows correspond to the transitions due to mobility.

It is worth noting that, for mobile users, the traffic conservation equation (2) no longer applies: the traffic arriving in region i is not equal in general to the traffic served in region i. Thus the mean throughput in region i is given by (1) but not by (3). Now the overall traffic conservation equation (5) still applies (in the absence of inter-cell mobility) so that the mean throughput in the cell is given by (6). It follows from (1) and (4) that

$$\gamma = \sum_i p'_i \gamma_i.$$

B. Round-robin policy

Under the round-robin policy, there is no longer an explicit expression for the stationary distribution of the Markov process $X(t)$, except in the limiting regime of infinite mobility where $\nu_{i,i+1}, \nu_{i+1,i} \to \infty$ for all $i < N$. In this case, the N queues are equivalent to one processor-sharing queue with service rate:
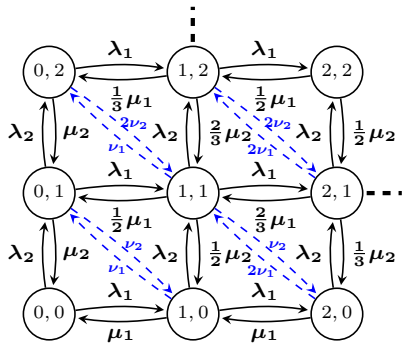
$$\bar{\mu} = \sum_i q_i \mu_i.$$

Fig. 4. State diagram for two regions under the round-robin policy.

We denote by $\bar{\rho} = \lambda/\bar{\mu}$ the corresponding load. The stationary distribution of $X(t)$ becomes:

$$\pi(x) = (1 - \bar{\rho})\bar{\rho}^n \binom{n}{x_1, \ldots, x_N} q_1^{x_1} \ldots q_N^{x_N},$$

with $n = \sum_i x_i$, under the stability condition $\bar{\rho} < 1$. The flow throughput is independent of the region where the flow starts and is given by:

$$\bar{\gamma} = \bar{\mu}\sigma (1 - \bar{\rho}).$$

The stability condition, which can be written

$$\rho < \frac{\bar{\mu}}{\mu}, \tag{9}$$

is actually independent of the mobility rates: at maximum load, flows tend to be infinitely long so that users move in the whole cell during their file transfers and have the mean physical rate $\bar{\mu}\sigma$. In the particular case where $p_i = q_i$ for all i (that is, the flow arrivals are distributed according to the distribution of the user positions), then $\bar{\mu} \geq \mu$ (the weighted arithmetic mean is larger than the corresponding weighted harmonic mean) and the maximum load increases, as observed in [4].
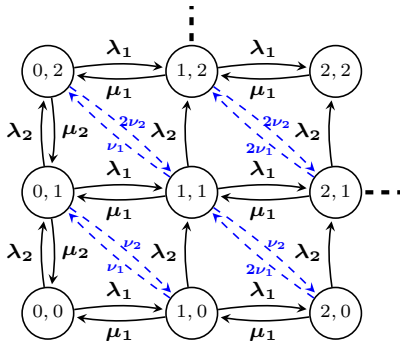
C. Max C/I policy

Fig. 5. State diagram for two regions under the max C/I policy.

Under the max C/I policy, it is also possible to derive the stationary distribution of the Markov process $X(t)$ in the limiting regime of infinite mobility. Given n active users, the probability that the user(s) with the best radio conditions is (are) in region i is given by

$$q_i(n) = \left(\sum_{j \geq i} q_j \right)^n - \left(\sum_{j > i} q_j \right)^n.$$

Since the max C/I policy only serves this (these) user(s), the mean service rate becomes:

$$\bar{\mu}(n) = \sum_i q_i(n)\mu_i.$$

Observe that $\bar{\mu}(n)$ is increasing, with

$$\bar{\mu}(1) = \bar{\mu} \quad \text{and} \quad \lim_{n \to \infty} \bar{\mu}(n) = \mu_1.$$

The stationary distribution of $X(t)$ is:

$$\pi(x) = \pi(0) \prod_{k=1}^n \frac{\lambda}{\bar{\mu}(k)} \times \binom{n}{x_1, \ldots, x_N} q_1^{x_1} \ldots q_N^{x_N},$$

with $n = \sum_i x_i$, under the stability condition $\lambda < \mu_1$.

This stability condition is in fact independent of the mobility rates. The maximum load is:

$$\rho < \frac{\mu_1}{\mu}, \tag{10}$$

which is always larger than in the absence of mobility.

D. Throughput in light traffic

The performance in light traffic (that is, when $\rho \to 0$) is independent of the scheduling policy. Indeed, a user when alone in the system is always allocated all radio resources. The mean throughput in region i is then equal to $\mu_i\sigma$. For the mean throughput in the cell, we consider the simple case of two regions. The probability that a user in region 1 moves to region 2 before leaving the system is:

$$\alpha_1 = \frac{\nu_1}{\nu_1 + \mu_1}.$$

Similarly, the probability that a user in region 2 moves to region 1 before leaving the system is:

$$\alpha_2 = \frac{\nu_2}{\nu_2 + \mu_2}.$$

We deduce the mean duration of a flow initiated in region 1

$$\frac{1}{\nu_1 + \mu_1} + \alpha_1 \frac{1}{\nu_2 + \mu_2} + \alpha_1\alpha_2 \frac{1}{\nu_1 + \mu_1} + \ldots$$

$$= \sum_{n \geq 0} (\alpha_1\alpha_2)^n \frac{1}{\nu_1 + \mu_1} + \alpha_1 \sum_{n \geq 0} (\alpha_1\alpha_2)^n \frac{1}{\nu_2 + \mu_2}$$

$$= \frac{1}{1 - \alpha_1\alpha_2} \left(\frac{1}{\nu_1 + \mu_1} + \frac{\alpha_1}{\nu_2 + \mu_2} \right).$$

Similarly, the mean duration of a flow initiated in region 2 is

$$\frac{1}{1 - \alpha_1\alpha_2} \left(\frac{1}{\nu_2 + \mu_2} + \frac{\alpha_2}{\nu_1 + \mu_1} \right).$$

By (6), the cell-average throughput in light traffic is

$$\gamma = \frac{\nu_2\mu_1 + \nu_1\mu_2 + \mu_1\mu_2}{\nu_1 + \nu_2 + p_1\mu_2 + p_2\mu_1}\sigma. \tag{11}$$

Observe that this throughput depends on user's mobility: $\gamma = \mu\sigma$ when $\nu_1, \nu_2 \to 0$ (no mobility) while $\gamma \to \bar{\mu}\sigma$ when $\nu_1, \nu_2 \to +\infty$ (infinite mobility).

E. Numerical results

Consider two regions with $p_1 = p_2 = 1/2$, $\mu_1 = 8$, $\mu_2 = 2$ and $\sigma = 1$. The mobility rates are symmetric, that is $\nu_1 = \nu_2 = \nu$. We consider the static case where $\nu = 0$, a case with moderate mobility $\nu = 0.1$, and the case of very high (nearly infinite) mobility, $\nu = 100$. The results are obtained by the numerical evaluation of the stationary distribution of the Markov process $X(t)$ and shown in Figure 6.
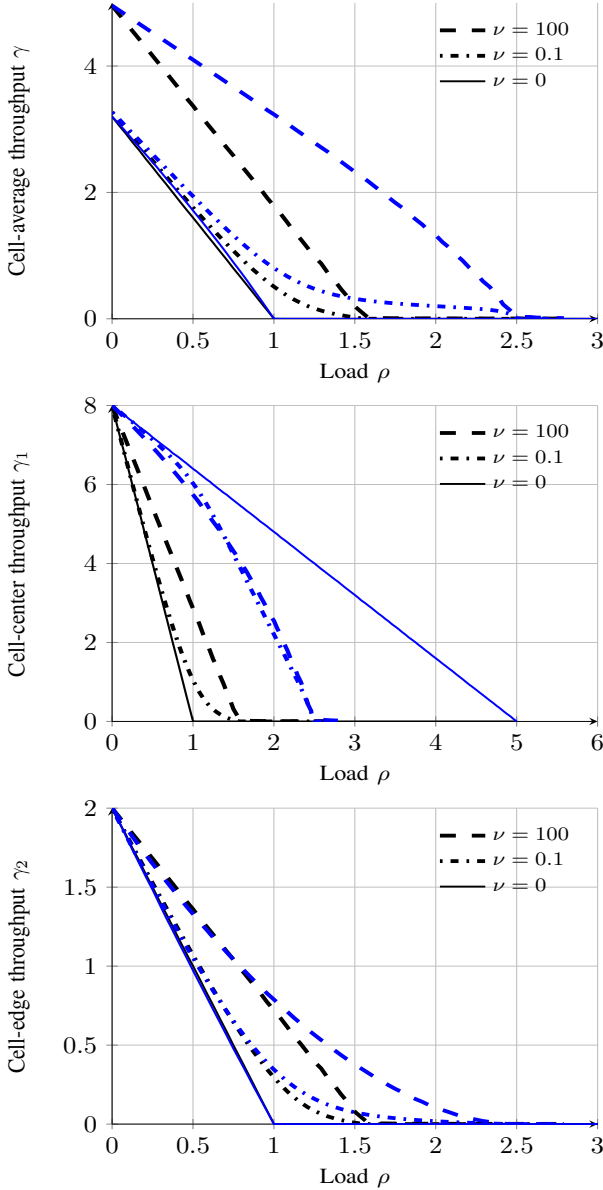

Fig. 6. Throughput performance for two regions with and without mobility, under round-robin (black) and max C/I (blue).

As expected, the maximum load when users move is independent of the mobility rate ν; it is approximately equal to 1.6 under the round-robin policy, in accordance with (9), and equal to 2.5 under the max C/I policy, in accordance with (10), both corresponding to huge gains compared to the static case. The cell-average throughput in light traffic is independent of the scheduling policy but very sensitive to the mobility rate, as predicted by (11).

We observe that mobility improves the cell-average throughput for both scheduling policies, with a higher gain under the max C/I policy. However, mobility may decrease the cell-center throughput (depending on the load) since users may then suffer from bad radio conditions. Surprisingly, the max C/I policy outperforms the round-robin policy for cell-edge users in case of mobility, even in case of moderate mobility $\nu = 0.1$: the fact that cell-edge users are not scheduled (in the presence of center-cell users) is compensated by the fact that cell-center users complete more rapidly their file transfers.

IV. MULTI-CLASS EXTENSION

We have assumed so far that all users have the same mobility behavior. In particular, they are all static or all mobile. We here extend the results to multiple classes of mobility, where each user may be static or mobile, as illustrated by Figure 7.
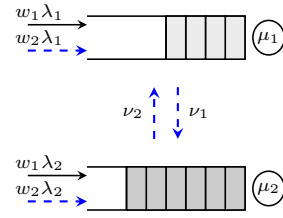

Fig. 7. Queuing model for two regions and two classes of mobility.

A. Mobility model

We consider K classes of mobility. For instance, each class represents a range of speeds (pedestrian, train, etc). Class-1 users are static. For any $k > 1$, a class-k user moves from region i to region $i + 1$ and from region $i + 1$ to region i (for $i < N$) at respective rates $\nu_{i,i+1}^{(k)}$ and $\nu_{i+1,i}^{(k)}$. We denote by $w^{(k)}$ the probability that a new flow is initiated by a class-k user. Let $X_i^{(k)}(t)$ be the number of active class-k flows in region i at time t. The Markov process $X(t)$ now has dimension $N \times K$. The probability that a class-k user is in region i satisfies

$$q_i^{(k)} \propto \prod_{j=1}^{i-1} \frac{\nu_{j,j+1}^{(k)}}{\nu_{j+1,j}^{(k)}}.$$

B. Round-robin policy

In the limiting regime where mobile users have infinite speeds, class-k users are served at rate

$$\bar{\mu}^{(k)} = \sum_i q_i^{(k)} \mu_i,$$

for any $k > 1$, while class-1 users are static and thus served at the mean rate μ. Let $\bar{\rho}^{(1)} = w^{(1)} \lambda / \mu$ and $\bar{\rho}^{(k)} = w^{(k)} \lambda / \bar{\mu}^{(k)}$ for $k > 1$. The queuing system corresponds to a multi-class processor-sharing queue of load $\bar{\rho} = \sum_k \bar{\rho}^{(k)}$. We deduce the stationary distribution of the Markov process $X(t)$:

$$\pi(x) = (1 - \bar{\rho}) \bar{\rho}^n \times \binom{\sum_i x_i^{(1)}}{x_1^{(1)}, \ldots, x_N^{(1)}} \left(\frac{p_1 \mu}{\mu_1}\right)^{x_1^{(1)}} \ldots \left(\frac{p_N \mu}{\mu_N}\right)^{x_N^{(1)}}$$

$$\times \prod_{k>1} \binom{\sum_i x_i^{(k)}}{x_1^{(k)}, \ldots, x_N^{(k)}} q_1^{(k)x_1^{(k)}} \ldots q_N^{(k)x_N^{(k)}},$$

with $n = \sum_{i,k} x_i^{(k)}$, under the stability condition $\bar{\rho} < 1$, that is

$$\rho < \frac{1}{w^{(1)} + \sum_{k>1} w^{(k)} \mu/\bar{\mu}^{(k)}}. \tag{12}$$

The mean throughput of static users in region i is

$$\gamma_i^{(1)} = \mu_i \sigma (1 - \bar{\rho}),$$

while the mean throughput of mobile users depends on their class $k > 1$:

$$\gamma^{(k)} = \bar{\mu}^{(k)} \sigma (1 - \bar{\rho}).$$

The cell-average throughput is

$$\gamma = \bar{\mu} \sigma (1 - \bar{\rho}),$$

where

$$\bar{\mu} = \frac{1}{w^{(1)}/\mu + \sum_{k>1} w^{(k)}/\bar{\mu}^{(k)}}.$$

C. Max C/I policy

These results can hardly be extended to the max C/I policy. Only the stability condition is explicit. At maximum load, mobile users are served only in region 1; we deduce the stability conditions for static users in region i:

$$\lambda \left(w^{(1)} \sum_{j \leq i} \frac{p_j}{\mu_j} + (1 - w^{(1)}) \frac{1}{\mu_1} \right) < 1,$$

that is

$$\rho < \frac{1}{w^{(1)} \sum_{j \leq i} \frac{p_j \mu}{\mu_j} + (1 - w^{(1)}) \frac{\mu}{\mu_1}}. \tag{13}$$

The overall stability condition is:

$$\rho < \frac{1}{w^{(1)} + (1 - w^{(1)}) \frac{\mu}{\mu_1}}. \tag{14}$$

D. Throughput in light traffic

When $\rho \to 0$, the mean throughput of static users in region i is $\mu_i \sigma$ while the mean throughput of mobile users of class $k > 1$ is, for two regions,

$$\gamma^{(k)} = \frac{\nu_2^{(k)} \mu_1 + \nu_1^{(k)} \mu_2 + \mu_1 \mu_2}{\nu_1^{(k)} + \nu_2^{(k)} + p_1 \mu_2 + p_2 \mu_1} \sigma.$$

E. Numerical Results

Figure 8 gives the results for the same example as in §III-E but with two classes of users: a static class and a mobile class, with $w^{(1)} = w^{(2)} = 50\%$. The maximum load is approximately 1.2 for the round-robin policy, as given by (12), and 1.4 under the max C/I policy, as given by (14), while static cell-center users have positive throughput up to load 3.3, in accordance with (13). We observe that, even for *static* users at the cell edge, the max C/I policy outperforms the round-robin policy in the presence of mobile users.
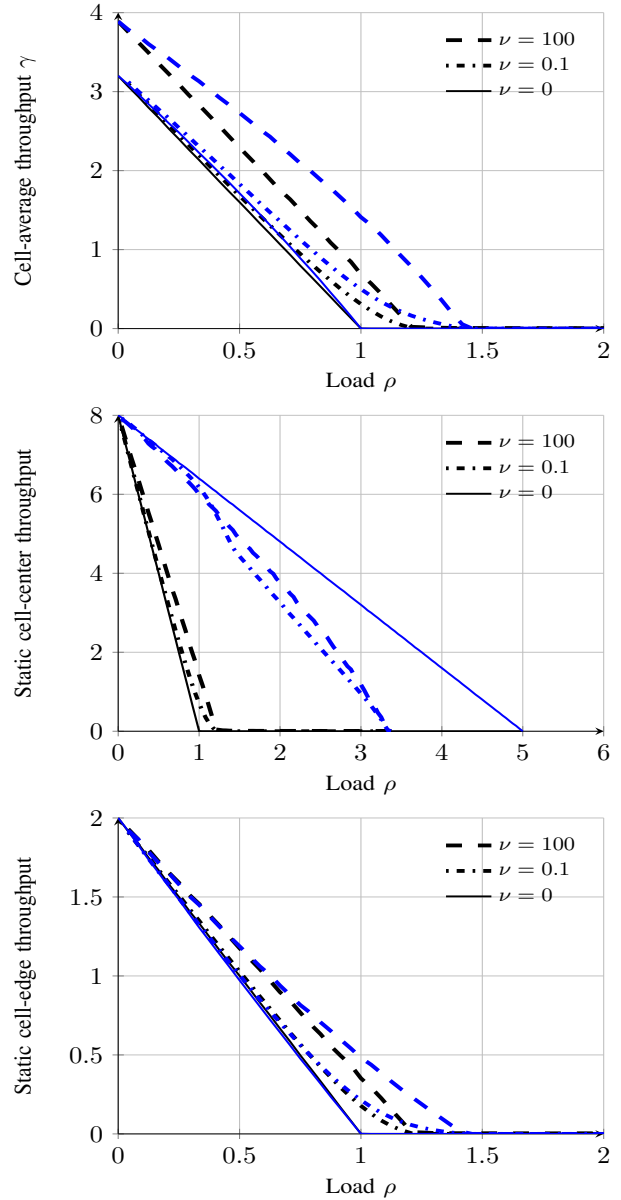

Fig. 8. Throughput performance for two regions with two classes of mobility, under round-robin (black) and max C/I (blue).

V. IMPACT OF INTER-CELL MOBILITY

We now take hand-overs into account: users may leave the cell during the file transfer. Similarly, some users may arrive from neighboring cells. We assume that handovers occur only in region N (the cell edge) and that the incoming and outgoing handover rates are equal.

Let ν_h be the outgoing handover rate: each user in region N tends to leave the cell at rate ν_h. To balance incoming and outgoing handovers, we assume that incoming handovers are generated according to a Poisson process of intensity $E(X_N)\nu_h$. Note that the net flow arrival rate in region N is still equal to λ_N. In particular, the traffic conservation equation (5) applies and the stability condition remains the same as without handovers for both policies.

These properties suggest that inter-cell mobility has a

limited impact on throughput performance, in the presence of intra-cell mobility. This is confirmed by the results of Figure 9, obtained for the example of §III-E with a handover rate $\nu_h = \nu = 0.1$, using fixed-point iterations to estimate the incoming handover rate.
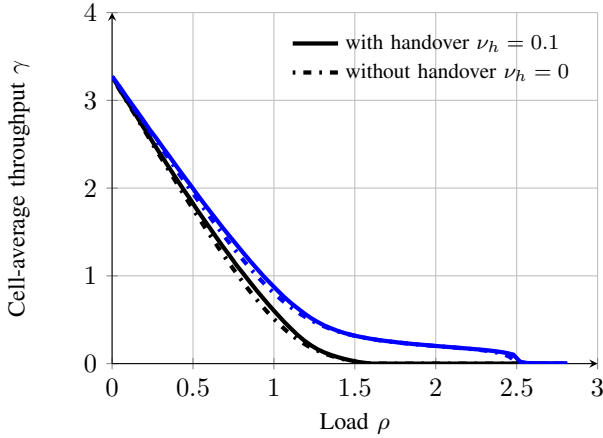

Fig. 9. Throughput performance for two regions with both intra-cell and inter-cell mobility, under round-robin (black) and max C/I (blue).

VI. IMPACT OF FAST FADING

We have so far ignored the impact of fast fading, which is the essence of opportunistic schedulers [16]. Fast fading occurs on a time scale of the order of ms, much shorter than the typical time scale of slow fading, of the order of second. This is why these two phenomena can be considered separately.

Consider n active users in region i, without any active users in regions $1, \ldots, i-1$ so that only these n users are scheduled. Assuming Rayleigh fading, users' SINR (signal-to-interference-plus-noise ratios) $S_1, S_2 ... S_n$ are exponentially distributed. The resulting SINR gain is given by:

$$F(n) = E\left[\max(S_1, S_2, \ldots, S_n)\right] = 1 + \frac{1}{2} + \ldots + \frac{1}{n}.$$

The throughput gain is typically lower due to the concavity of the data rate with respect to SINR. Assuming that the physical data rate is proportional to $\log_2(1+S)$ for SINR S, we deduce the rate gain of opportunistic scheduling:

$$G(n) = \frac{E[\log_2(1 + \max(S_1, S_2, \ldots, S_n))]}{E[\log_2(1 + S_1)]}.$$

The corresponding gains are shown in Figure 10.

In any state x such that $x_i > 0$, the total throughput of users in region i becomes:

$$\mu_i \sigma \phi_i(x) G(x_i).$$

The results obtained for the same scenario as in §III-E (without inter-cell mobility) are shown in Figure 11. Compared to Figure 6, the max C/I policy brings a significant throughput gain due to fast fading, in addition to that due to slow fading.
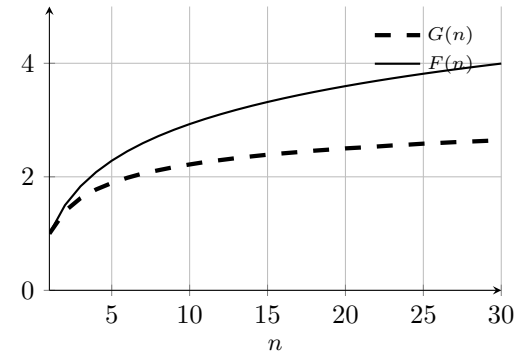

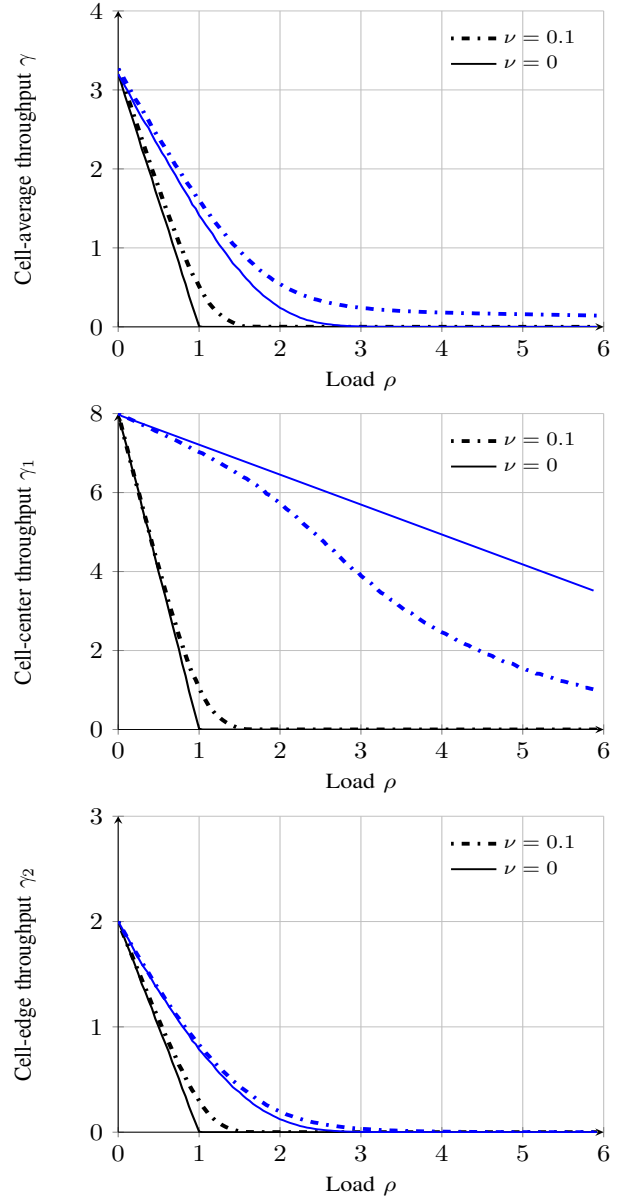Fig. 10. SINR gain $F(n)$ and rate gain $G(n)$ for n scheduled users.

Fig. 11. Throughput performance for two regions with fast fading, under round-robin (black) and max C/I (blue).

VII. System simulations

A. Simulation setting

We now validate the results of previous sections by system-level simulations based on the LTE technology. We consider the 21 hexagonal cells formed by 7 tri-sector sites (a reference site surrounded by 6 interferring sites). Both eNodeB and mobile devices are equipped with 2 antennas. Spatially correlated shadowing maps for different sites are generated each $3s$, with time correlation. Each TTI (of 1 ms), 4 spatially correlated fading complex values are generated for each eNodeB to mobile device pair. The main parameters are summarized in Table I.

Network topology	Macro cells only
Environment	Urban
Context	outdoor
Inter-site distance	500 m
Resolution	5 m
PathLoss	ITU Model
shadowing std	4 dB
shadowing cross correlation (between sites)	0.5
shadowing correlation distance	30 m
Radio access technology	LTE
Number of tx/rx antennas	MIMO (2x2)
number of streams	1
Receiver	MRC
Codebook	3GPP
carrier frequency	2 GHz
Bandwidth	10 MHz

TABLE I.   Simulation setting.

Traffic consists of file transfers only. Flows arrive according to a Poisson process with uniform spatial distribution. File sizes are generated from an exponential distribution with mean $\sigma = 0.5$ MB. Scheduling decisions are taken in each eNodeB at each TTI. In order to evaluate the proposed model, the two previous schedulers are implemented. The round-robin scheduler serves users in a cycle way, independently of their radio conditions; the max C/I scheduler chooses at each TTI the user with the highest instantaneous SINR.

We study two scenarios, with and without mobility. With mobility, each user chooses a direction at random and moves along this direction at a constant speed of 30, 60, 100 and 300 km/h until the file transfer is completed. The network is a torus in this case to avoid border effects (that is, a user leaving cell 20 from the south arrives in cell 6, 7 or 13 to the north according to the direction). We simulate 10 minutes and estimate for each scenario the cell-average throughput as the ratio of mean flow size to mean flow duration.

B. Results

Figure 12 shows the normalized cell-average throughput (in bit/s/Hz) as a function of the arrival rate per eNodeB in each scenario. Results are very close to those obtained by analysis.

VIII. Conclusion

We have quantified the gains induced by mobility in cellular data networks. Both the round-robin and max C/I policies improve throughput performance, with a proportionally higher gain for the latter. By prioritizing cell-center users, the max C/I gives the opportunity to mobile cell-edge users to be served in better radio conditions.
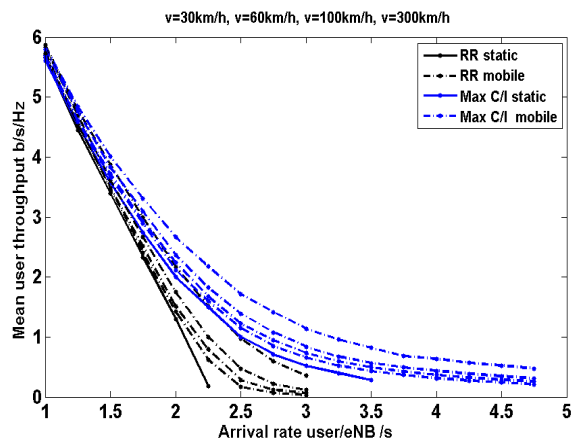

Fig. 12.   Throughput performance obtained by system-level simulation.

Although the focus of this paper is on the downlink, we think that similar performance gains can be obtained on the uplink. Other interesting problems for future research include opportunistic scheduling policies that prioritize short flows, as considered in [1] for instance, as well as heterogeneous networks consisting of both picocells and macro cells [12].

References

[1] S., Aalto, A. Penttinen, P. Lassila, P. Osti. On the optimal trade-off between SRPT and opportunistic scheduling. In: *ACM SIGMETRICS 2011*.

[2] U. Ayesta, M. Erausquin, P. Jacko (2010). A modeling framework for optimizing the flow-level scheduling with time-varying channels. Performance Evaluation, 67(11).

[3] T. Bonald. A score-based opportunistic scheduler for fading radio channels. In: *Proc. European Wireless 2004*.

[4] T. Bonald, S.C. Borst and A. Proutière (2004). How mobility impacts the flow-level performance of wireless data systems. In: *Proc. IEEE Infocom 2004*.

[5] T. Bonald, S.C. Borst, N. Hedge, M. Jonckheere and A. Proutière (2009). Flow-level performance and capacity of wireless networks with user mobility. Queuing Systems Theory Applications, 63(1-4), 131-164.

[6] T. Bonald, M. Feuillet, Network performance analysis. Wiley ed., 2011.

[7] S. Borst, A. Proutière and N. Hedge. Capacity of wireless data networks with intra- and inter-cell mobility. In: *Proc. IEEE Infocom 2006*.

[8] S. Borst, A. Proutière and N. Hedge. Mobility-driven scheduling in wireless networks. In: *Proc. IEEE Infocom 2009*.

[9] S.E. Elayoubi, T. Chahed and G. Hébuterne (2007). Mobility-aware admission control scheme in the downlink of third-generation wireless systems. IEEE Trans. on Vehicular Technology, 56(1), 245-259.

[10] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. In: *IEEE Infocom 2001*.

[11] M. K. Karray (2011). User's mobility effect on the performance of wireless cellular networks serving elastic traffic. Wireless Net., 17(1).

[12] V. Kavitha, S. Ramanath, E. Altman (2011). Spatial queueing for analysis, design and dimensioning of Picocell networks with mobile users. Performance Evaluation, 68(8), 710-727.

[13] J. Kim, B. Kim, J. Kim, Y. H. Bae (2013). Stability of flow-level scheduling with Markovian time-varying channels. Performance Evaluation, 70(2), 148-159.

[14] L. Kleinrock, Queueing Systems. Volume 1: Theory. Wiley, 1975.

[15] T.E. Kolding, K.I. Pedersen, J. Wigard, F. Frederiksen and P.E. Mogensen (2003). High speed downlink packet access: WCDMA evolution. IEEE Vehicular Technology Society News, 50(1), 4-10.

[16] D. Tse. Multiuser Diversity in Wireless Networks. In: *Wireless Communication seminar, Stanford University Apr. 2001*.