Flow-Level Performance of Intra-site Coordination in Cellular Networks

Ahlem Khlass
Orange Labs & Telecom ParisTech
Paris, France
ahlem.khlass@telecom-paristech.fr

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

Salah Eddine Elayoubi
Orange Labs
Issy-Les-Moulineaux, France
salaheddine.elayoubi@orange.com

***Abstract*—In this paper, we assess the ability of intra-site coordination schemes to combat inter-cell interference in cellular networks. We first focus on the static scheme proposed by the 3GPP standards where coordination is always performed in hand-over regions. Through the analysis of a flow-level model, we show that this scheme indeed improves cell-edge throughput at low loads but may make the system unstable at high loads, due to the suboptimal allocation of radio resources. We then propose a dynamic scheme where coordination decisions depend on the loads of the different sectors. Results show that this dynamic scheme behaves like the static one at low loads, outperforms it at medium loads and preserves the stability region of the network at high loads.**

***Index Terms* – Cellular networks, sector coordination, flow-level modeling, queuing theory.**

I. Introduction

Inter-cell interference is a major issue in cellular networks. It does not only make users at cell edge suffer from low throughputs but also decrease the overall network capacity since these users consume a significant proportion of the radio resources [1]. Cell coordination has been proposed as an efficient way for reducing inter-cell interference by allowing several base stations to transmit data simultaneously to the same user [2]. It has been recently adopted by 3GPP standards for Beyond 3G. Release 11 introduced cell coordination for High Speed Packet Access (HSPA) in the Multipoint transmission feature [3]. In LTE-Advanced, discussions are undergoing for including cell coordination in Release 12 under the name of Coordinated MultiPoint (CoMP) [4].

There are two broad classes of cell coordination schemes: inter-site coordination, where cells of different sites are coordinated, and intra-site coordination, where cells (or sectors) of the same site are coordinated [5]. Inter-site coordination remains difficult to implement for several reasons. First, it requires a high-speed, low-delay backhaul network so that base stations can exchange control signals in nearly real time and thus take pertinent coordination decisions [3]. Second, a tight synchronization has to be performed between cells as coherent transmissions are usually required to reduce interference. Third, clustering base stations into groups of coordinated cells is needed in order to reduce the complexity of the scheduling schemes [6]. Intra-site coordination is a much simpler and more practical scheme since all scheduling decisions are local. It is the focus of the present paper.

The major challenge of cell coordination is to control the strong coupling of scheduling decisions in the different cells. Indeed, scheduling a user in a given cell may require to silent some other cells, leading to the following dilemma:

1) Cell coordination increases substantially the data rates of users when scheduled.
2) Cell coordination leads to situations where only few uses are simultaneously scheduled in the network, thus decreasing the proportion of time where individual users are scheduled.

This tradeoff between higher achievable rates (due to lower interference) and higher radio resource consumption (due to joint transmissions to the same user) is arguably the key issue when designing a coordination scheme. We analyse this tradeoff in the practical interesting case of intra-site coordination in HSDPA networks, which carry the majority of the data traffic of many mobile operators.

The first attempts to evaluate the performance of coordination schemes (see e.g. [2] and references therein) focused on physical layer aspects and were based on so-called full buffer simulations, as proposed by 3GPP [7]. In this approach, a fixed number of users is simulated in each cell with a nearly complete emulation of the physical and MAC layers. The advantage of these simulations is that they are highly accurate with respect to the lower layers as they take into account a complete channel model including path loss, shadowing and fast fading. However, they do not capture the dynamic aspect of traffic, which is critical for coordination schemes (see point 2) of the dilemma above). Due to this weakness, so-called finite buffer simulators have been recently proposed (see the FTP simulation model in [7]). These simulations add the dynamic traffic layer whereby short-lived data flows are generated at random times, mimicking the user behavior. Such simulations have recently been used in [5] and [8] to evaluate the performance of coordination schemes in LTE-Advanced and HSPA networks, respectively. Substantial throughput gains are observed, especially at the cell edge.

Although finite buffer simulations provide useful insights into the efficiency of coordination schemes, they are computationally intensive, with typical simulation times in the range of days. Analytical methods are practically interesting and, to the best of our knowledge, have not yet been proposed in the

considered dynamic setting with flow arrivals and departures. Indeed, the underlying model corresponds to a set of coupled queues, which is known to be intractable beyond some very specific cases [9]. Bounds and approximations on the impact of inter-cell interference on user throughput have been derived in [10] and [11], respectively, but without coordination scheme.

The present paper is a first attempt for evaluating the performance of cell coordination schemes analytically. We focus on intra-site coordination and start by investigating the static coordination scheme as proposed by the 3GPP standard [3]. We show in particular that this scheme improves cell-edge throughput at low loads but may make the system unstable at high loads, due to the suboptimal allocation of radio resources. Therefore, we propose an advanced coordination scheme which takes into account the load of the different sectors. Results show that this dynamic scheme behaves like the static one at low loads, outperforms it at medium loads and preserves the stability region of the network at high loads.

The remainder of this paper is organized as follows. We first describe the considered coordination schemes. We then present the model, including both radio and traffic aspects. Sections IV and V are dedicated to the performance analysis of the static and dynamic coordination schemes, respectively. Section VI concludes the paper.

II. Coordination Schemes

A. Multipoint transmission

Recently, the concept of Multipoint Transmission has been introduced in 3GPP Release 11 for cell coordination in HSDPA networks [3]. Two mechanisms are standardized: *single-point transmission*, whereby the user equipment (UE) is allowed to be attached to two base stations (BSs) but can be served only by one of them at each time slot, depending on its instantaneous channel quality, and *multi-flow transmission*, whereby two BSs can simultaneously schedule independent transport blocks to the UE. We focus on the latter, which is expected to provide the highest throughput gains. Since we consider intra-site coordination only, the BSs subject to coordination serve different sectors of the same site.

B. Static and dynamic schemes

In the absence of coordination, each UE is associated to the so-called serving BS which offers the best signal. When the pilot signal from a non-serving BS exceeds that from the serving BS by some threshold, say δ, a handover (HO) mechanism is triggered in order to select the best BS to switch to. When coordination is performed, the non-serving BS with the best signal coordinates with the serving BS in order to schedule independent transport blocks to the UE. Note that only those UEs in the HO region, as defined by the power threshold parameter δ, benefit from coordination.

We consider two coordination schemes:

Static scheme: This is the coordination scheme adopted in 3GPP Release 11 [3]. We refer to this scheme as static since a UE is simultaneously scheduled by two BSs serving adjacent sectors whenever it lies in the HO region of these sectors. This scheme proves beneficial at low loads only, as revealed by the analysis of section IV; at high loads, the coordination brings an additional load which maintains the network in congestion due to the lack of available resources.

Dynamic scheme: This scheme aims at providing coordination while preserving the radio resources. Specifically, a UE in the HO region of two sectors receives data from the non-serving BS only if the latter has no UE to serve in its sector. We refer to this scheme as dynamic since the coordination decisions now depend on the state of the sectors and vary with time. The condition imposed on the coordination decisions may look too restrictive at first sight, since only empty sectors coordinate with other sectors. Our results of section V indicate that this is not the case: the dynamic scheme activates coordination and improves the cell-edge throughput at low loads, since sectors are often empty, and blocks coordination at high loads, so as to preserve scheduling resources and prevent the network from congestion.

C. Scheduling algorithm

Both the static scheme and the dynamic scheme require a centralized scheduler to decide which UEs to serve at each time-slot. This is not a strong requirement in the context of intra-site coordination considered here. Although the scheduling algorithm should in practice be opportunistic and exploit multi-user diversity in order to improve the spectral efficiency of the system, we consider a simple algorithm where UEs are selected at random, as specified below. This is a good approximation of the resource allocation achieved by any blind, fair algorithm like round-robin and by any opportunistic algorithm when radio channel variations are limited or too fast to be exploited.

At each timeslot, the algorithm works as follows:

Data: Set A of active sectors (with UEs to be served)
Result: Set U of UEs to schedule
while $A \neq \emptyset$ **do**
  pick up at random a UE u served by a sector s in A;
  $U \leftarrow U \cup \{u\}$;
  $A \leftarrow A \setminus \{s\}$;
  if coordination **then**
    $c \leftarrow$ coordinating sector;
    $A \leftarrow A \setminus \{c\}$;
  end
end

The coordination decision depends on the scheme (static or dynamic) and on the location of user u (in the HO region or not). For the dynamic scheme, there is coordination only if the interfering sector is empty, so that $c \notin A$: the update $A \leftarrow A \setminus \{c\}$ is useless in this case.

III. Model

A. Radio aspects

Consider a site consisting of K sectors indexed by k. Each sector is served by a dedicated BS. In the absence of coor-

dination, the signal-to-noise-plus-interference ratio (SINR) of UE u served by BS s is given by:

$$\mathrm{SINR}_{u,s} = \frac{P_{u,s}}{\sum_{k \neq s} P_{u,k} + I_u + N},$$

where $P_{u,k}$ denotes the power received by UE u from BS k, I_u is the inter-site interference received by UE u and N is the thermal noise. The resulting data rate is $f(\mathrm{SINR}_{u,s})$ for some function f; a standard model for f is provided by the Shannon capacity of the Gaussian channel $f(s) = W \log(1+s)$, where W is the channel bandwidth. This only provides an upper bound, however, and more realistic values derived from real systems are used in §III-B below.

When coordination is performed, UE u receives simultaneously two different transport blocks from the serving BS s and from the coordinating BS c, leading to two flows with respective SINRs:

$$\mathrm{SINR}'_{u,s} = \frac{P_{u,s}}{\sum_{k \neq s,c} P_{u,k} + I_u + N}$$

and

$$\mathrm{SINR}'_{u,c} = \frac{P_{u,c}}{\sum_{k \neq s,c} P_{u,k} + I_u + N}.$$

Note that, as required by 3GPP, the mobile is equipped with a 3i-receiver capable of perfectly canceling the interference from the coordinated BSs [3]. The data rate is the sum $f(\mathrm{SINR}'_{u,s}) + f(\mathrm{SINR}'_{u,c})$, higher than in the absence of coordination.

B. Achievable rates

For the numerical applications, we consider the homogeneous tri-sectorized network of Figure 1. Figure 2 gives the achievable rates in some reference cell with respect to the distance to the BS and the angle relative to the main direction of the antenna, with and without coordination, under the following assumptions:

- the radio channel is PA3 in a dense urban area with an inter-site distance of 500 meters, see [12] for more details,
- inter-cell interference is calculated based on two interfering rings of sites, corresponding to a total of 19 sites,
- the rate function f, which takes fast fading into account, is that given in [13],
- in the presence of coordination, the HO region is assumed to represent a fraction α of the sector area.

We observe that the achievable rates are almost doubled when coordination is applied. Note that this does not take the impact of traffic into account, however. It is precisely the aim of this paper to analyze the trade-off between higher achievable gains and lower scheduling opportunities due to the resource consumption of the coordination schemes.

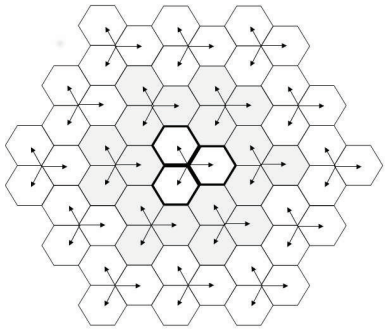

Fig. 1. A tri-sectorized hexagonal network.

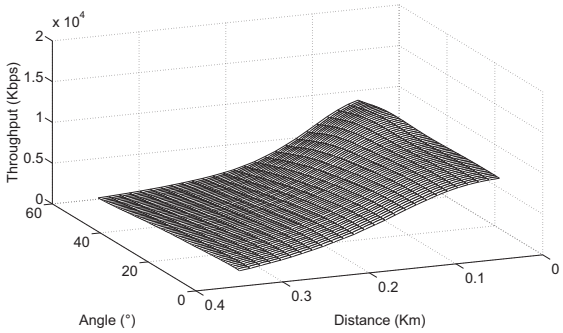

(a) Without coordination.

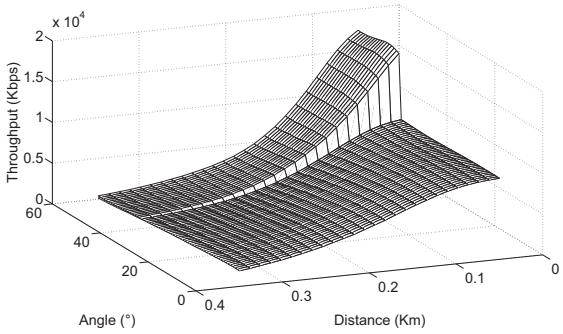

(b) With coordination ($\alpha = 30\%$).

Fig. 2. Achievable rates in a tri-sectorized hexagonal network.

C. Traffic characteristics

To analyze this trade-off, we need a simple model for the random behavior of users. We assume that data flows arrive according to a Poisson process of intensity λ_k in each sector k of the reference site. Note that this assumption is reasonable since users typically behave independently. Each flow stays in the system as long as the corresponding data have not been successfully transmitted to the UE. Flow sizes are assumed to be independent and exponentially distributed with mean σ bits. Most results are in fact approximately insensitive to the flow size distribution beyond the mean, as explained below. The traffic intensity in sector k is $\lambda_k \sigma$ in bps.

As seen above, the achievable rate of a UE depends on

its location in the sector. We consider an arbitrary set S of classes of UEs indexed by i to reflect the different radio conditions in the considered site. All UEs of the same class have approximately the same location and thus the same serving BS and the same data rate when scheduled. We denote by S_k the set of classes whose serving BS is k; UEs of class i have rate R_i when scheduled without coordination and rate R'_i when scheduled with coordination. The corresponding flows are referred to as class-i flows in the following. A flow arriving in sector k is of class $i \in S_k$ with probability p_i, with:

$$\sum_{i \in S_k} p_i = 1. \tag{1}$$

Note that class-i flows arrive according to a Poisson process with intensity[1] $\lambda_i \equiv \lambda_k p_i$ in sector k, contributing to the traffic intensity $\lambda_i \sigma$ (in bps).

D. Throughput metric

Let $X_i(t)$ be the number of class-i flows at time t. The vector $X(t) = (X_i)_{i \in S}$ is an irreducible Markov process whose transition rates depend on the considered coordination scheme, as described in the following two sections. Assuming that this Markov process is ergodic, the mean duration τ_i of class-i flows follows from Little's law:

$$\tau_i = \frac{E(X_i)}{\lambda_i}.$$

Defining the flow throughput (in bits) as the ratio of the mean flow size (in bits) to the mean flow duration (in s), we deduce the class-i flow throughput:

$$\gamma_i = \frac{\lambda_i \sigma}{E(X_i)}. \tag{2}$$

This is the ratio of the traffic intensity of class i (in bps) to the mean number of class-i flows. This throughput metric reflects user experience, accounting both for the radio conditions, through class i, and for the random nature of traffic, through the stationary distribution of the Markov process $X(t)$. The flow throughput averaged over the whole site is given by:

$$\bar{\gamma} = \frac{\sum_{i \in S} \lambda_i \sigma}{\sum_{i \in S} E(X_i)} = \frac{\sum_{i \in S} \lambda_i}{\sum_{i \in S} \lambda_i / \gamma_i}. \tag{3}$$

This is the harmonic mean of the per-class flow throughputs weighted by the corresponding traffic intensities.

IV. Static coordination scheme

In this section, we derive the performance of the static scheme; we start with the reference case without coordination, which also corresponds to the limiting case where there is no HO region.

[1]This slight abuse of notation whereby the parameter is defined by its subscript is used throughout the paper; it makes mathematical expressions much easier to read.

A. No coordination

Consider class-i flows in sector k. Since the mean flow size is σ, the mean scheduling time required at BS k to complete a class-i is σ / R_i. Since class-i flows arrive at rate λ_i, the corresponding load is:

$$\rho_i = \frac{\lambda_i \sigma}{R_i}. \tag{4}$$

At the flow time-scale, slots are very short (in the range of ms) so that the random, fair scheduling algorithm described in §II-C can be viewed as a processor-sharing service discipline per sector. Specifically, the number of flows in sector k evolves like the number of customers in a multi-class processor-sharing queue with load:

$$\rho_k = \sum_{i \in S_k} \rho_i. \tag{5}$$

Note that, in view of (1) and (4), we have:

$$\rho_k = \frac{\lambda_k \sigma}{R_k},$$

where R_k is the harmonic mean data rate over sector k:

$$R_k = \frac{1}{\sum_{i \in S_k} \frac{p_i}{R_i}}.$$

Under the stability condition $\rho_1 < 1, \ldots, \rho_K < 1$ (no sector is in overload), the stationary distribution of the state x describing the number of flows of each class in the site is given by:

$$\pi(x) = \prod_{k=1}^{K} (1 - \rho_k) \frac{x_k!}{\prod_{i \in S_k} x_i!} \prod_{i \in S_k} \rho_i^{x_i}. \tag{6}$$

We deduce the mean number of class-i flows in sector k:

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

and from (2) and (4), the flow throughput of any class $i \in S_k$:

$$\gamma_i = R_i (1 - \rho_k). \tag{7}$$

Thus the flow throughput is equal to the maximum data rate R_i when $\rho_k = 0$ (no traffic) and decreases linearly with the sector load ρ_k.

B. Coordinating two sectors

To gain insight into the behavior of the static coordination scheme, we start with the simple case of $K = 2$ sectors. We refer to zone 12 as the HO region of sectors 1 and 2. We denote by $C_{12} \subset S$ the set of classes in zone 12; UEs of class $i \in C_{12}$ are served simultaneously by BS 1 and 2 at rate R'_i. The sets $C_1 = S_1 \setminus C_{12}$ and $C_2 = S_2 \setminus C_{12}$ define zones 1 and 2 of sectors 1 and 2, respectively, where UEs are not subject to coordination; UEs of class $i \in C_j$ are served by BS j at rate R_i, for $j = 1, 2$. Let $Z_j(t)$ be the total number of flows in zone j at time t, for $j = 1, 2, 12$. Denote by $Z(t)$ the corresponding vector. Whenever $Z(t) = z$, the considered

scheduler selects UEs in zones 1 and 2 (assuming $z_1, z_2 > 0$) a fraction of time:

$$\bar{\phi}_1(z) = \bar{\phi}_2(z) = \frac{z_1 + z_2}{z_1 + z_2 + z_{12}}, \tag{8}$$

and UEs in zone 12 a fraction of time:

$$\bar{\phi}_{12}(z) = \frac{z_{12}}{z_1 + z_2 + z_{12}}. \tag{9}$$

Indeed, since UEs in zone 12 require the simultaneous transmission from both BSs, the algorithm described in §II-C cannot select a UE in zone 12 if the first UE u selected at random is in zone 1 or 2. Since time-sharing is fair in each zone, the system corresponds to a set of three coupled processor-sharing queues with state-dependent service rates given by (8)-(9) and respective loads:

$$\bar{\rho}_1 = \frac{\bar{\lambda}_1 \sigma}{\bar{R}_1}, \quad \bar{\rho}_2 = \frac{\bar{\lambda}_2 \sigma}{\bar{R}_2}, \quad \bar{\rho}_{12} = \frac{\bar{\lambda}_{12} \sigma}{\bar{R}'_{12}}, \tag{10}$$

with arrival rates:

$$\bar{\lambda}_1 = \sum_{i \in C_1} \lambda_i, \quad \bar{\lambda}_2 = \sum_{i \in C_2} \lambda_i, \quad \bar{\lambda}_{12} = \sum_{i \in C_{12}} \lambda_i,$$

and harmonic mean data rates:

$$\bar{R}_1 = \frac{\sum_{i \in C_1} p_i}{\sum_{i \in C_1} \frac{p_i}{R_i}}, \quad \bar{R}_2 = \frac{\sum_{i \in C_2} p_i}{\sum_{i \in C_2} \frac{p_i}{R_i}},$$

and

$$\bar{R}'_{12} = \frac{\sum_{i \in C_{12}} p_i}{\sum_{i \in C_{12}} \frac{p_i}{R'_i}}.$$

Note that, unlike in zones 1 and 2, the data rates in zone 12 benefit from the coordination gains.

The service rates (8)-(9) satisfy the balance property so that the corresponding system is a Whittle queueing network [14]. In particular, the stationary distribution of the stochastic process $Z(t)$ is given by:

$$\pi(z) = \frac{(1 - \bar{\rho}_1 - \bar{\rho}_{12})(1 - \bar{\rho}_2 - \bar{\rho}_{12})}{(1 - \bar{\rho}_{12})} \times \binom{z_1 + z_2 + z_{12}}{z_1 + z_2} \bar{\rho}_1^{z_1} \bar{\rho}_2^{z_2} \bar{\rho}_{12}^{z_{12}}, \tag{11}$$

under the stability condition:

$$\bar{\rho}_1 + \bar{\rho}_{12} < 1, \quad \bar{\rho}_2 + \bar{\rho}_{12} < 1. \tag{12}$$

Moreover, the results are *insensitive* to the flow size distribution beyond the mean: it is not necessary to assume an exponential flow size distribution.

We deduce from the stationary distribution (11) the mean number of flows in each zone:

$$E(Z_1) = \frac{\bar{\rho}_1}{1 - \bar{\rho}_1 - \bar{\rho}_{12}},$$

$$E(Z_2) = \frac{\bar{\rho}_2}{1 - \bar{\rho}_2 - \bar{\rho}_{12}},$$

$$E(Z_{12}) = \frac{\bar{\rho}_1}{1 - \bar{\rho}_{12}} \left(1 + \frac{\bar{\rho}_1}{1 - \bar{\rho}_1 - \bar{\rho}_{12}} + \frac{\bar{\rho}_2}{1 - \bar{\rho}_2 - \bar{\rho}_{12}}\right).$$

and from (2), the flow throughput in each zone:

$$\bar{\gamma}_1 = \bar{R}_1(1 - \bar{\rho}_1 - \bar{\rho}_{12}), \quad \bar{\gamma}_2 = \bar{R}_2(1 - \bar{\rho}_2 - \bar{\rho}_{12}), \tag{13}$$

$$\bar{\gamma}_{12} = \frac{\bar{R}'_{12}(1 - \bar{\rho}_{12})}{1 + \frac{\bar{\rho}_1}{1 - \bar{\rho}_1 - \bar{\rho}_{12}} + \frac{\bar{\rho}_2}{1 - \bar{\rho}_2 - \bar{\rho}_{12}}}. \tag{14}$$

The mean number of flows of class i being proportional to the load of this class in the corresponding queue, we also obtain the flow throughput of each class within each zone:

$$\forall i \in C_1, \quad \gamma_i = R_i(1 - \bar{\rho}_1 - \bar{\rho}_{12}),$$

$$\forall i \in C_2, \quad \gamma_i = R_i(1 - \bar{\rho}_2 - \bar{\rho}_{12}),$$

$$\forall i \in C_{12}, \quad \gamma_i = \frac{R'_i(1 - \bar{\rho}_{12})}{1 + \frac{\bar{\rho}_1}{1 - \bar{\rho}_1 - \bar{\rho}_{12}} + \frac{\bar{\rho}_2}{1 - \bar{\rho}_2 - \bar{\rho}_{12}}}.$$

Comparing these expressions to (7) provides the first insights into the trade-off between achievable data rates and radio resource consumption mentioned in the introduction. When $\bar{\rho}_1, \bar{\rho}_2, \bar{\rho}_{12} \to 0$, we get $\gamma_i \to R'_i$ in the coordination zone: there is a throughput gain of R'_i/R_i for class i, as expected. When load increases, the coordination scheme may be detrimental. In the case for two symmetric sectors for instance, the stability condition (12) is stricter than the usual condition $\rho_k < 1$ for each sector k whenever $\bar{R}'_{12} < 2\bar{R}_{12}$, where $\bar{R}_{12}$ denotes the mean data rate in the HO region in the absence of coordination:

$$\bar{R}_{12} = \frac{\sum_{i \in C_{12}} p_i}{\sum_{i \in C_{12}} \frac{p_i}{R_i}}.$$

Thus the ability of the static scheme to improve performance at high load critically depends on the ratio $\beta = \bar{R}'_{12}/\bar{R}_{12}$, which we refer to as the coordination gain β in the following. We shall see that typical values of β are less than 2 so that the static scheme in fact reduces the stability region and degrades performance at high load.

C. Coordinating three sectors

We now consider the practically interesting case of $K = 3$ sectors. The analysis can be readily extended to any number of sectors. As previously, we refer to zone $k, k+1$ as the HO region between sectors k and $k+1$, for all $k = 1, 2, 3$ (with modulo K notations) and denote by $C_{k,k+1} \subset S_k \cap S_{k+1}$ the corresponding classes. Zone k refers to the non-HO region of sector k; it corresponds to classes $S_k \setminus (C_{k,k-1} \cup C_{k,k+1})$.

Let $Z(t)$ be the vector of the number of flows in each zone. When $Z(t) = z$, the considered scheduler selects UEs in zone 1 a fraction of time:

$$\bar{\phi}_1(z) = \frac{z_1 + z_{23}}{|z|} + \frac{z_2}{|z|} \frac{z_1 + z_3}{z_1 + z_3 + z_{31}} + \frac{z_3}{|z|} \frac{z_1 + z_2}{z_1 + z_2 + z_{12}} \tag{15}$$

and UEs in zone 12 a fraction of time:

$$\bar{\phi}_{12}(z) = \frac{z_{12}}{|z|} + \frac{z_3}{|z|} \frac{z_{12}}{z_1 + z_2 + z_{12}}, \tag{16}$$

with $|z| = z_1 + z_2 + z_3 + z_{12} + z_{23} + z_{13}$. We obtain (16) for instance by observing that a UE in zone 12 is either selected first by the scheduling algorithm, with probability $z_{12}/|z|$, or after the selection of a UE in zone 3, with probability $z_3/|z|$; in the latter case, only UEs in zones 1, 2 or 12 can be selected in parallel. Expression (15) can be derived in the same way, and the timeslot allocation in the other zones follow by symmetry.

Since time-sharing is fair in each zone, the system corresponds to a set of six coupled processor-sharing queues with state-dependent service rates (15)–(16) and loads defined by (10), as in the case of two sectors. Unfortunately, the service rates violate the balance property so that the system is no longer a Whittle network, see [14].

To get explicit expressions for the flow throughput, we proceed by approximation which we validate numerically. Specifically, we decouple the six zones by neglecting those which are not directly linked by coordination. For instance, we focus the study on zones 1, 2, 12 and neglect the impact of sector 3, beyond the load induced on BS 1 and 2. In doing so, we reduce the analysis to the case of two sectors considered above. In particular, the stationary distribution of the number of flows in zones $1, 2, 12$ is approximated by (11), where the loads of zones 1 and 2 are replaced by $\bar\rho_1 + \bar\rho_{13}$ and $\bar\rho_2 + \bar\rho_{23}$, respectively. We deduce the following approximations for the flow throughputs in zones 1 and 12:

$$\bar\gamma_1 \approx \bar R_1(1 - \bar\rho_1 - \bar\rho_{12} - \bar\rho_{13}), \tag{17}$$

$$\bar\gamma_{12} \approx \frac{\bar R'_{12}(1 - \bar\rho_{12})}{1 + \dfrac{\bar\rho_1 + \bar\rho_{13}}{1 - \bar\rho_1 - \bar\rho_{12} - \bar\rho_{13}} + \dfrac{\bar\rho_2 + \bar\rho_{23}}{1 - \bar\rho_2 - \bar\rho_{12} - \bar\rho_{23}}}, \tag{18}$$

under the stability condition:

$$\bar\rho_1 + \bar\rho_{12} + \bar\rho_{13} < 1,$$
$$\bar\rho_2 + \bar\rho_{21} + \bar\rho_{23} < 1,$$
$$\bar\rho_3 + \bar\rho_{31} + \bar\rho_{32} < 1.$$

Importantly, the analysis is consistent in the sense that the expression for the flow throughput in zone 1 is the same whether the approximation is applied to zones $1, 2, 12$ or to zones $1, 3, 13$. Per-class flow throughputs follow by replacing mean achievable data rates by per-class achievable data rates.

In order to validate the analytical approximation, we consider an isotropic site with three symmetric sectors. The system is then defined by the following three parameters:

- α, the fraction of traffic in the coordination zone,
- η, the ratio of mean achievable data rates in and out of the HO region, in the absence of coordination,
- β, the coordination gain, that is the ratio of mean achievable data rates in the HO region with and without coordination.

We have:

$$\alpha = \frac{\bar\lambda_{12}}{\bar\lambda_1 + \bar\lambda_{12}}, \quad \eta = \frac{\bar R_{12}}{\bar R_1}, \quad \beta = \frac{\bar R'_{12}}{\bar R_{12}}.$$

The normalized flow throughput averaged over the whole site is shown in Figure 3 for $\alpha = 0.3$, $\eta = 0.5$ and $\beta = 1.5$. The chosen throughput unit is the mean achievable data rate in the absence of coordination; in view of (3) and (7), the flow throughput is then equal to $1 - \rho$ in the absence of coordination, where ρ is the load of each sector. The analytical expression derived from (17) and (18) is compared with the numerical solution of the Markov process describing the state of the queuing system, assuming an exponential flow size distribution. We observe that the approximation is very accurate. Other simulations, not reported here, show that the results are practically insensitive to the flow size distribution, which is not surprising given the insensitivity property of the system for $K = 2$ sectors.

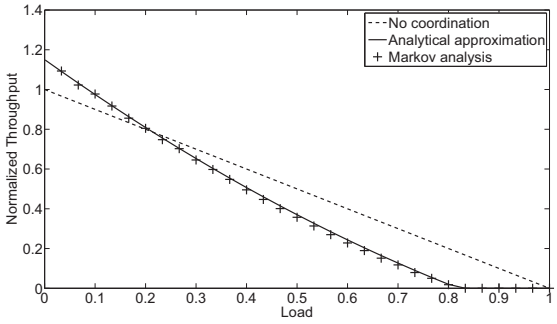

Fig. 3. Validation of the approximation (static coordination scheme).

Using the same approximation with the achievable data rates of §III-B, we obtain the flow throughputs of Figure 4 for different sizes α of the coordination zone. The flow throughput is averaged over the site (corresponding to $\bar\gamma$) and at cell edge (corresponding 5% of users with the lowest throughput). We observe a significant performance improvement at low load (up to 25% for the cell edge, 20% on average) at the expense of a reduction of the stability region. Indeed, the coordination gain β derived from the achievable data rates of §III-B is typically equal to 1.5, that is less than the critical value $\beta = 2$ where the stability region remains unchanged. Static coordination is only beneficial at low load.

V. Dynamic coordination scheme

The dynamic scheme consists in allowing a BS to coordinate and serve a UE not in its sector only if it has no other UE to serve. We shall see that this simple scheme is sufficient to improve performance at any load.

A. Coordinating two sectors

Like for the static scheme, we start with the case $K = 2$. We use the same notation, except that we now differentiate between zone 12, consisting of classes in the HO region served by BS 1, and zone 21, consisting of classes in the HO region served by BS 2. We denote by $Y_1 = Z_1 + Z_{12}$ and $Y_2 = Z_2 + Z_{21}$ the total number of flows in sectors 1 and 2, respectively. The system is more complex than in the static case since the two processor-sharing queues associated with each sector now interact at the class level. We decouple the

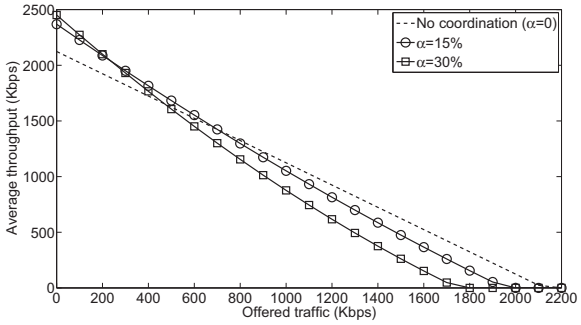

(a) Average throughput.

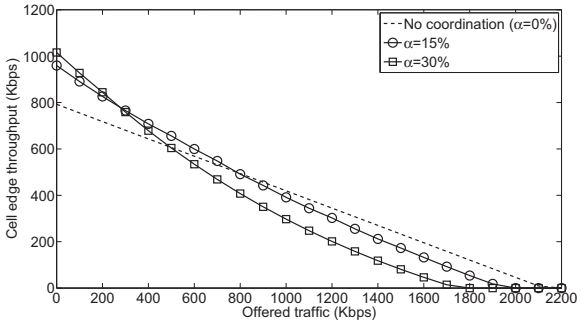

(b) Cell edge throughput.

Fig. 4. Flow throughput for different sizes of the coordination zone.

system by replacing the state-dependent coordination gain of each class by its mean, so that the coordination gain in zones 12 and 21 are respectively given by:

$$\bar{\beta}_{12} = P(Y_2 > 0) + \frac{\bar{R}'_{12}}{\bar{R}_{12}} P(Y_2 = 0),$$

$$\bar{\beta}_{21} = P(Y_1 > 0) + \frac{\bar{R}'_{21}}{\bar{R}_{21}} P(Y_1 = 0).$$

Sectors 1 and 2 now behave as two independent multi-class processor-sharing queue with respective loads:

$$\rho_1 = \frac{\bar{\lambda}_1 \sigma}{\bar{R}_1} + \frac{\bar{\lambda}_{12}\sigma}{\bar{\beta}_{12}\bar{R}_{12}}, \quad \rho_2 = \frac{\bar{\lambda}_2 \sigma}{\bar{R}_2} + \frac{\bar{\lambda}_{21}\sigma}{\bar{\beta}_{21}\bar{R}_{21}}.$$

In particular, we have:

$$P(Y_1 > 0) = \rho_1 \quad \text{and} \quad P(Y_2 > 0) = \rho_2,$$

from which we deduce $\bar{\beta}_{12}$ and $\bar{\beta}_{21}$. The flow throughputs in sector 1 are then given by:

$$\forall i \in C_1, \quad \gamma_i = R_i(1 - \rho_1),$$

$$\forall i \in C_{12}, \quad \gamma_i = (R_i \rho_2 + R'_i(1 - \rho_2))(1 - \rho_1).$$

The expressions for sector 2 follow by symmetry.

It is worth observing that the stability region is preserved: since $\bar{\beta}_{12}, \bar{\beta}_{21} > 1$, the stability condition $\rho_1 < 1, \rho_2 < 1$ is not stricter than that in the absence of coordination. Both stability regions in fact coincide since $\bar{\beta}_{12}, \bar{\beta}_{21} \to 1$ at high load.

B. Coordinating three sectors

The approach is similar for K sectors. We take $K = 3$ to simplify the notation. The coordination gain in zone 12 is given by:

$$\bar{\beta}_{12} = P(Y_2 > 0) + \frac{\bar{R}'_{12}}{\bar{R}_{12}} P(Y_2 = 0) + \frac{1}{2}\left(1 - \frac{\bar{R}'_{12}}{\bar{R}_{12}}\right) P(Y_2 = 0, Y_3 > 0) E(\frac{Z_{32}}{Y_3} | Y_3 > 0).$$

Indeed, a UE in zone 12 benefits from coordination only if sector 2 is empty and no UE in zone 32 benefits from coordination. When sector 2 is empty and UEs in zones 12 and 32 are scheduled, we assume that BS 2 selects BS 1 or 3 uniformly at random for coordination, hence the factor $\frac{1}{2}$. Sectors 1, 2, 3 behave as three independent multi-class processor-sharing queue with respective loads:

$$\rho_1 = \frac{\bar{\lambda}_1 \sigma}{\bar{R}_1} + \frac{\bar{\lambda}_{12}\sigma}{\bar{\beta}_{12}\bar{R}_{12}} + \frac{\bar{\lambda}_{13}\sigma}{\bar{\beta}_{13}\bar{R}_{13}},$$

$$\rho_2 = \frac{\bar{\lambda}_2 \sigma}{\bar{R}_2} + \frac{\bar{\lambda}_{21}\sigma}{\bar{\beta}_{21}\bar{R}_{21}} + \frac{\bar{\lambda}_{23}\sigma}{\bar{\beta}_{23}\bar{R}_{23}},$$

$$\rho_3 = \frac{\bar{\lambda}_3 \sigma}{\bar{R}_3} + \frac{\bar{\lambda}_{31}\sigma}{\bar{\beta}_{31}\bar{R}_{31}} + \frac{\bar{\lambda}_{32}\sigma}{\bar{\beta}_{32}\bar{R}_{32}}.$$

In particular, we have:

$$P(Y_2 > 0) = \rho_2, \ P(Y_3 > 0) = \rho_3, \ E(\frac{Z_{32}}{Y_3} | Y_3 > 0) = \frac{\bar{\rho}_{32}}{\rho_3},$$

with:

$$\bar{\rho}_{32} = \frac{\bar{\lambda}_{32}\sigma}{\bar{\beta}_{32}\bar{R}_{32}},$$

from which we deduce the coordination gains $\bar{\beta}_j$ for each coordination zone $j = 12, 21, 23, 32, 13, 31$. The stability condition $\rho_1 < 1$, $\rho_2 < 1$, $\rho_3 < 1$ is the same as without coordination and the flow throughputs in zones 1 and 12 are given by:

$$\forall i \in C_1, \quad \gamma_i = R_i(1 - \rho_1),$$

$$\forall i \in C_{12}, \quad \gamma_i = (R_i \rho_2 + R'_i(1 - \rho_2) + \frac{R_i - R'_i}{2}(1 - \rho_2)\bar{\rho}_{32})(1 - \rho_1).$$

The expressions for the other zones follow by symmetry.

In the symmetric case, with the parameters α, η, β introduced in §IV-C, the common coordination gain $\bar{\beta}$ satisfies the equation:

$$\bar{\beta} = \rho + \beta(1 - \rho) + \frac{1}{4}(1 - \beta)\rho(1 - \rho)\frac{\alpha}{\alpha + (1 - \alpha)\eta\bar{\beta}}, \quad (19)$$

where ρ denotes the common sector load, given for a flow arrival rate λ per sector by:

$$\rho = \lambda(1 - \alpha + \frac{\alpha}{\eta\bar{\beta}}). \quad (20)$$

The stability condition is $\rho < 1$ and the flow throughputs in zone 1 and zone 12 are respectively given by:

$$\bar{\gamma}_1 = \bar{R}_1(1 - \rho) \quad \text{and} \quad \bar{\gamma}_{12} = \bar{\beta}\bar{R}_{12}(1 - \rho).$$

The normalized flow throughput averaged over the whole site is shown in Figure 5 for $\alpha = 0.3$, $\eta = 0.5$ and $\beta = 1.5$. As in Figure 3, the throughput unit is the mean achievable data rate in the absence of coordination. The flow throughput derived from the solution of the fixed-point equation (19), (20) is compared with the numerical solution of the Markov process describing the state of the queuing system, assuming an exponential flow size distribution. Again, the approximation is very accurate and performance is practically insensitive ot the flow size distribution.

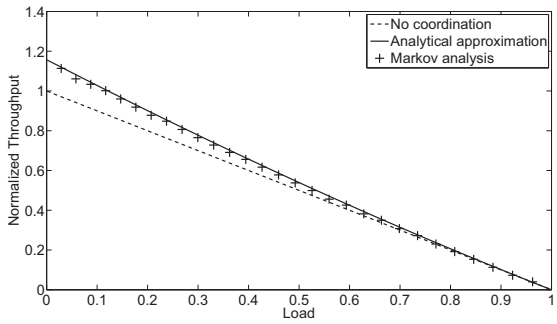

Fig. 5. Validation of the approximation (dynamic coordination scheme).

Figures 6 shows the flow throughputs averaged over the site and at cell edge, respectively, using the achievable data rates of §III-B. The first observation is that, in contrast with the static scheme, the dynamic scheme always improves performance. Indeed, while the dynamic scheme is able to achieve the same gains as the static scheme at low loads (up to 25% for the cell edge, 20% on average), it preserves the stability region and prevents throughput degradation at high loads.

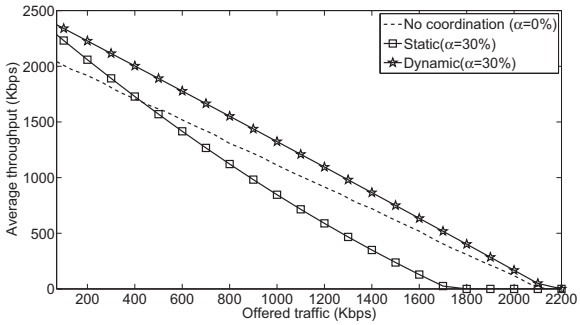

(a) Average throughput.

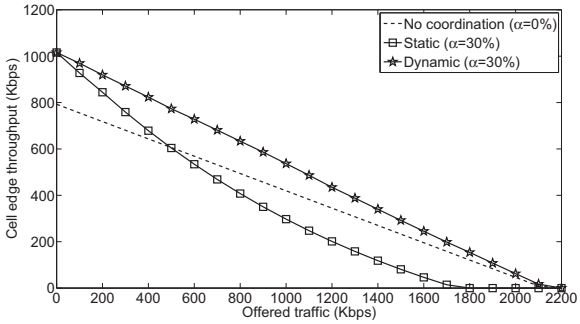

(b) Cell edge throughput.

Fig. 6. Flow throughput under the static and dynamic schemes.

VI. Conclusion

This paper analyses the performance of intra-site coordination in cellular data networks, accounting for the random, dynamic nature of traffic. Using a simple flow-level model, we have highlighted the shortcomings of the coordination scheme proposed in the 3GPP standards, we refer to as the *static* scheme: the stability region reduces due to the over-consumption of time-slots, which leads to a throughput degradation at high load, both on average in the site and at the cell edge. Thus we have proposed a *dynamic* scheme that activates coordination only when necessary. Results show a throughput improvement at any load.

The case of more advanced coordination and scheduling schemes exploiting multi-user diversity is left as an interesting track for further research.

References

[1] T. Bonald and A. Proutière, "Wireless downlink data channels: user performance and cell dimensioning," in *Proceedings of the 9th annual international conference on Mobile computing and networking*, ser. MobiCom '03. New York, NY, USA: ACM, 2003, pp. 339–352.

[2] M. Karakayali, G. Foschini, and R. Valenzuela, "Network coordination for spectrally efficient communications in cellular systems," *Wireless Communications, IEEE*, vol. 13, no. 4, pp. 56 –61, aug. 2006.

[3] *HSDPA Multipoint Transmission (Release 11)*, 3GPP TR 25.872 Std.

[4] *Requirements for further advancements for Evolved Universal Terrestrial Radio Access (E-UTRA) (LTE-Advanced)*, 3GPP TR 36.913 Std.

[5] D. Lee, H. Seo, B. Clerckx, E. Hardouin, D. Mazzarese, S. Nagata, and K. Sayana, "Coordinated multipoint transmission and reception in LTE-advanced: deployment scenarios and operational challenges," *IEEE Communications Magazine*, vol. 50, no. 2, pp. 148–155, 2012.

[6] F. Huang, Y. Wang, J. Geng, M. Wu, and D. Yang, "Clustering approach in Coordinated Multi-Point transmission/reception system," in *Vehicular Technology Conference Fall (VTC 2010-Fall), 2010 IEEE 72nd*, sept. 2010, pp. 1 –5.

[7] *Evolved Universal Terrestrial Radio Access (E-UTRA); Further advancements for E-UTRA physical layer aspects (Release 9)*, 3GPP TR 36.814 V9.0.0 (2010-03) Std.

[8] C. Zhang, Y. Chang, S. Liu, and D. Yang, "System-level analysis and evaluation of SF-DC transmit mode in HSPA system," in *Wireless Communications and Networking Conference (WCNC), 2012 IEEE*, april 2012, pp. 1185 –1190.

[9] G. Fayolle and R. Iasnogorodski, "Two coupled processors: The reduction to a Riemann-Hilbert problem," *Probability Theory and Related Fields*, vol. 47, pp. 325–351, 1979, 10.1007/BF00535168.

[10] T. Bonald, S. Borst, N. Hegde, and A. Proutière, "Wireless data performance in multi-cell scenarios," in *Proceedings of ACM SIGMETRICS*, 2004, pp. 378–387.

[11] A. J. Fehske and G. P. Fettweis, "Aggregation of variables in load models for interference-coupled cellular data networks," *IEEE ICC*, 2012.

[12] *Digital cellular telecommunications system (Phase 2+); Radio network planning aspects*, ETSI TR 143 030 Std.

[13] A. Saadani and J.-B. Landre, "Realistic performance of HSDPA evolution 64-QAM in macro-cell environment," in *Vehicular Technology Conference, 2009. VTC Spring 2009. IEEE 69th*, april 2009, pp. 1 –5.

[14] T. Bonald and A. Proutière, "A queueing analysis of data networks," in *Queueing Networks: A Fundamental Approach*, ser. International Series in Operations Research & Management Science, R. J. Boucherie and N. M. van Dijk, Eds. Springer Verlag, 2011, vol. 154, pp. 643–699.