A System-level Assessment of Uplink CoMP in LTE-A Heterogeneous Networks

Mohamad Tavakoli Sanij, Claudio Casetti,
Politecnico di Torino, Italy

Abstract—In LTE-Advanced networks, the steady demand for higher data rates by users is met through several techniques. One of the most promising solutions is CoMP (Coordinated MultiPoint), which allows the involvement of multiple eNB in the transmission and reception process with a marked increase in throughput for users at the edge of cells. While downlink CoMP has been the focus of many works in the literature, in this paper we examine *Uplink* CoMP with Coordinated Scheduling in Heterogeneous Networks running LTE-A. We investigate its system-level performance through simulation in various realistic scenarios with frequency-selective Rayleigh fading. Specifically, we are interested in comparing the performance of Uplink CoMP with different cell types as well as different user participation to CoMP transmissions. Our simulation results confirm that edge users benefit from Uplink CoMP, although the overall throughput decreases. In order to mitigate the latter effect, we introduce two new parameters, called CoMP margin and CoMP Pool Percentage (CPP), and provide guidelines for their effective use.

I. INTRODUCTION

LTE-Advanced meets the challenge raised by powerful, mobile devices and bandwidth-hungry applications by investing in solutions such as carrier aggregation, higher order MIMO, relay nodes and Coordinated Multipoint (CoMP) transmission/reception. The latter, in particular, is envisioned to be one of the most important techniques in LTE-Advanced to improve the throughput and functionality of cell borders [1], [2]. CoMP allows users to have multiple data transmission and reception from/toward multiple cooperating eNodeBs (eNBs), increasing the utilization factor of the network. Resource allocation in the uplink can leverage the availability of additional connection points where the signal from the User Equipment (UE) is processed, ultimately providing UEs with increased throughput. Additionally, a significant part of the interference caused by neighboring cells can be seen as a useful received signal thanks to CoMP, provided those cells are part of the Coordinated Reception Point (CRP) set. This is especially important in critical regions, in terms of interference, like cell edges. Finally, in the case of joint multicell scheduling, CoMP introduces a reduction in the backhaul load by requiring only scheduling data to be transferred between coordinated eNBs.

Arguably, CoMP is most appealing in the uplink direction since it does not require UE modifications: indeed, users need not be aware that there is any kind of cooperation among receiving eNBs. UEs are merely scheduled for transmission on a set of frequencies that happens to be split among different eNBs, although they still retain standard signalling channels through only one of these eNBs, usually referred to as the *serving* cell. In this work we focus on uplink CoMP from a system point of view. Specifically, we are interested in comparing through simulation the performance of uplink CoMP in a heterogeneous scenario with different user participation to CoMP transmissions. Some works have already investigated uplink CoMP both in simulation and through field trials. In particular, while [3] has mainly addressed early field trial results in both downlink and uplink CoMP, [4], [5] has focused on an uplink CoMP testbed using commercial equipment, comparing the bitrate gain introduced by CoMP selection combining and soft combining. In [6], the authors have simulated a system implementing both uplink CoMP and ICIC (Inter-Cell Interference Coordination), showing that the latter has improved performance when combined with CoMP. From an analytical standpoint, [7] has introduced trade-offs between efficient backhaul usage and maximum CoMP gain, while [8] has investigated new uplink scheduling techniques for CoMP based on spatially uncorrelated UEs and aiming at maximizing the long-term throughput.

Our contribution confirms the findings of previous works as far as the throughput gain for edge users is concerned, but introduces two novel observations that can spur future investigations on CoMP systems and lead to the design of new resource allocation algorithms:

- We look at four different scenarios where we restrict the type of cells that can be in the CRP set, identifying pros and cons of each of them.
- We introduce a parameter called CoMP Pool Percentage (CPP), which quantifies the fraction of PRBs[1], that are reserved for UEs using a specific eNB as CRP (out of the resources nominally available to that eNB). Our simulations show that the setting of CPP must be carefully gauged depending on the number of CoMP users and the scenario.

The paper is organized as follows. In Sec. II, we outline the details of the LTE-A system model that we are considering. The four CoMP scenarios are introduced in Sec III, while, in Sec. IV, we present the concept of CoMP Pool Percentage. Sec. V contains an in-depth discussion of simulation results. Finally, Sec. VI concludes the paper.

---

[1]In OFDMA, users are allocated a specific number of subcarriers for a limited amount of time: they are referred to as Physical Resource Blocks (PRBs) in LTE standards

II. LTE-A System Model

In this section, we carefully describe the system model we used in our simulations, highlighting the necessary *simplifications* that we have chosen to introduce in order to efficiently simulate a large area with 1,000 user roaming across it.

A. Radio coverage

We consider a 25-km^2 area A with a fixed number of macro cells, M, equal to 15 and a number of small cells, S, equal to 50. Each macro eNB has a three-sectorial antenna, irradiating a macro cell in each sector with a power $P_{TX} = 40dBm$. The antenna irradiation pattern determines a gain for antenna m expressed as follows:

$$G_m(\phi) = -\min\left[12\left(\frac{\phi}{\phi_{3dB}}\right), \hat{G}\right] \quad (1)$$

where:

$$\phi_{3dB} = 70°, -180° < \phi < 180°, \ \hat{G} = 25dB$$

As for small cells, each is irradiated by a single eNB at $P_{TX} = 30dBm$ in an omni-directional pattern with gain $G_s = 0_{dB}$. An open-loop power control regime is assumed for all cells. UE transmission power is assumed to be equally distributed across the allocated user bandwidth.

In order to better cover the area, the location of eNBs is such that they are placed as wide apart as possible, minimizing the overlap and the uncovered portion of the area.

The area is assumed to be populated by 1,000 UEs moving according to a simplified random direction model. The speed is randomly chosen for every UE according to a truncated Gaussian distribution (to avoid infinite speeds). The associated mean value of the Gaussian distribution is set to 3 km/h (slow UEs) and 50 km/h (fast UEs), with a standard deviation of 1.5 and 20 Km/h. The distribution of UE speed in our simulations is depicted in Figure 1. In order to maximize the effectiveness of small cell coverage, a *virtual hotspot* is defined within a 200m radius of each small cell. In each simulation, 60% of slow UEs are placed within a virtual hotspot. We point out that fast users do not have access to COMP functionalities, due to the large set up time it needs and that would likely result in those UEs leaving the cell before the multipoint connection is active. Fast users can however communicate with their serving cell and contribute with background traffic.

B. Path loss

Following [10] we compute the average channel path loss between a macro eNb antenna m and UE i at a distance $d_{m,i}$, as follows:

$$L_{m,i} = 128.1 + 37.6 \cdot log(d_{m,i}). \quad (2)$$

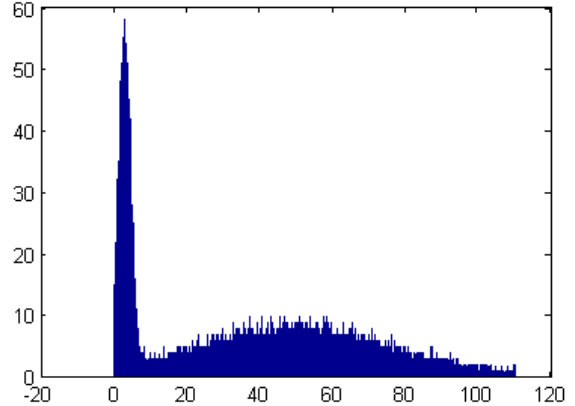Likewise, the average path loss between a small-cell eNb antenna s and UE i at a distance $d_{s,i}$ is:

Fig. 1. Distribution of UE Speed in Km/h

$$L_{s,i} = 140.7 + 36.7 \cdot log(d_{s,i}). \quad (3)$$

Since we focus on the uplink, we have considered the transmit power at the mobile device and the antenna gain at the receiver (eNB).

C. Fast fading

We model fast fading effects due to multipath using a frequency-selective Rayleigh model. Since the average distance between mobile users and eNBs is less than 150 meter, a multipath pattern with attenuation and five different delay taps, bounded at 110 ns, are set as an input for the Rayleigh fading process. Values of delay and attenuation are taken from the Extended Pedestrian A model (EPA) [9]. Although in the original model the number of taps is equal to seven, for the sake of simulation time and average distance, we decide to use just first five taps. Considering our scenario featuring high-speed and low-speed users (CoMP being used mainly by the latter), the maximum value of the doppler spread for CoMP users can be computed as:

$$\Delta F_{dopp} = \frac{V}{\lambda} \quad (4)$$

where V and λ are the relative mobile user velocity and central carrier frequency wavelength, respectively. With the numerical values used in our scenario, ΔF_{dopp} is less than 100 Hz. We thus assume that the fading process is approximately constant over a 10-ms coherence time. In order to simplify the simulation, we compute the fading losses every 10 ms instead of the TTI value of 1ms (Block Fading). A further simplification we introduce is the computation of fading loss from only 10 randomly selected PRBs, out of 100 PRBs. Thus, the expression of fading loss used in the computation of the Reference Signal Received Power (RSRP) between user i and n_{th} eNB is computed as:

$$A_{n,i} = \frac{\sum\limits_{PRB \in B} R_{loss}(n,i)}{10} \quad (5)$$

Parameter	Value
Number of small cells	50
Number of macro cells	15
Number of users	1000
Uplink Bandwidth	20 MHz
Handover threshold	3 dB
Max. Number of CRP	3
CRE Power Bias	6 dB
User Transmit Power	20 dBm
Number of PRBs	100
Number of Multipath	5
Fast fading Tap delays	0,30,70,90,110 ns
Fast fading attenuation	0, 1, 2, 3, 8 dB
Noise Power	-174 dBm/Hz

where $R_{loss}(n, i)$ is the Rayligh fast fading loss for user i and eNB number n and B is a set of ten randomly selected PRBs.

D. Serving Cell selection

In order to select the serving cell, the RSRP in downlink, defined in [11], is computed for every UE i from a macro eNB antenna m as:

$$RSRP(m, i) = P_{TX} + G_m(\phi) - L_{m,i} - A_{m,i} \quad (6)$$

while, from a small-cell eNB antenna s:

$$RSRP(s, i) = P_{TX} - L_{s,i} - A_{s,i} \quad (7)$$

Cell selection for every UE aims at maximizing the downlink RSRP, although a constant bias (6dB in our simulations) is added to the nominal small cell RSRP. This procedure, known as Cell Range Expansion (CRE), helps compensating for the power gap between macro and small cells, thus offloading UEs from the former to the latter. The CRE power bias is applied to both service cell selection as well as to the selection of Coordinated Reception Points for CoMP. In order to reduce the complexity of the simulation, the CRE bias at all small-cell eNBs are equal and static during the simulation. UEs on the move have their connection handed over to another cell when the RSRP (including CRE power bias) of the new cell is 3dB higher than the RSRP of the legacy cell. Such offset prevents ping-pong effects in the handover procedure.

The system parameters discussed above are summarized in Table I.

III. CoMP Scenarios

Uplink CoMP relies on measurements reported by the UE through standard signalling channels, comparing the received power of the candidate CoMP Reception Point (CRP) cell and that of the serving cell. If the RSRP of the CRP cell is higher than the RSRP of the serving cell minus a quantity called *CoMP margin* (measured in dB), then the candidate is selected as CRP. Intuitively, the higher the CoMP margin, the larger the number of UEs who satisfy the condition to become CoMP users. If more than one cell satisfies the

condition, the CRP are selected in increasing order of RSRP values, up to 3 CRP cells. Since back-haul conditions are not modeled in our simulation, the CoMP margin value does not depend on the network load, a feature that we will explore in future work. Likewise, the CoMP margin does not depend on cell sizes.

In our simulations, we consider four scenarios where macro and small cells can play different roles, as either serving cells or CRPs, or both. We acknowledge that each scenario entails the selection of different signalling overhead, which we however consider as negligible since we do not model the backhaul.

- *Intra:* CoMP is enabled just for one of the non-serving cells of the *same eNB* where the serving cell is located. In this case, the maximum number of CRP is limited to 2 and no small cell is involved as CRP. A small cell can be a serving cell in this scenario, but the user transmitting to it cannot use CoMP.
- *Inter:* CoMP is enabled for cells from any macro eNB, not restricted to cells of the serving eNB. Small cells cannot be used for CoMP.
- *Small:* although any cell can be a serving cell, only small cells can be CRP.
- *Het:* no restrictions in the roles of serving and CRP.

IV. Resource allocation and throughput computation

Our performance evaluation focuses on the allocation of radio resources among CoMP and non-CoMP users. At a system level, we assume that frequencies are allocated to eNBs out of a 20-MHz band according to the Fractional Frequency Planning scheme [12], with a frequency reuse factor of 4. We use a co-channel frequency deployment, hence small cells use the same frequency band that macro cells use. Our simulation model clearly accounts, through eq. (10), for any possible interference among users in different cells, who are allocated the same PRBs.

Within a cell, we simulate different resource splits. Every cell sets aside a fraction of resources, i.e., of Physical Resource Blocks (PRBs), for UEs that transmit to it as their serving cell. These resources are typically allocated to UEs close to the cell core. The remaining fraction of PRBs are reserved for UEs that use the cell as CRP and are thus close to the edge. We identify the latter fraction, i.e., the portion of PRBs that are reserved for CoMP as *CoMP Pool Percentage (CPP)*. In each simulation instance, the value of CPP is identical across all cells involved in CoMP (which, as we have seen, depends on the scenario). By definition, CPP is 0 for cells that do not participate to CoMP. We also point out that if the CPP of a cell is not completely allocated for lack of CoMP users, its PRBs are available to be scheduled for non-CoMP users.

Resources are then allocated to UE for their uplink communication based on the Proportional Fair (PF) scheduling policy [13]. The PF policy combines high throughput proportional fairness among all UEs by giving instantaneous priority to UEs with a high-quality channel rate and a

low average service rate. The user uplink throughput is computed from the number of PRBs that the scheduler allocates to each user, depending on the cell it communicates with (either serving or CRP). The per-PRB received power $P_{RX}(PRB)$ is computed by a macro eNB antenna m for UE i transmitting at power $P_{TX,i}$ as:

$$P_{RX}(PRB) = P_{TX,i} + G_m(\phi) - L_{m,i} - A_{m,i}(PRB) \quad (8)$$

while, by a small-cell eNb antenna s:

$$P_{RX}(PRB) = P_{TX,i} - L_{s,i} - A_{s,i}(PRB) \quad (9)$$

In the above equations $A_{m,i}(PRB)$ and $A_{s,i}(PRB)$ are losses introduced by the Rayleigh fast fading model for user i on the PRB allocated by a macro or small cell, respectively.

The throughput yielded by each PRB is derived from Shannon's formula, assuming a bandwidth of 180 kHz per PRB, a noise power density of -174 dBm/Hz, and a SINR computed as:

$$SINR = \frac{P_{RX}(PRB)}{P_{noise} + \sum\limits_{u \in S} P_{RX}(u)} \quad (10)$$

where P_{noise} is the noise power associated to the PRB bandwidth and S is a set of UEs, which are allocated the same PRB for their uplink communication.

UEs are assumed to be in saturation, i.e., they always have data traffic to upload and greedily use all the resources they are allocated by the eNB.

V. SIMULATION RESULTS

Simulations were run on a system-level simulator written in MATLAB and compatible with LTE specification. In the simulation, the throughput is averaged every second. Every simulation run lasts 60 seconds and results are averaged over 10 different runs. Such an amount of simulation time allows the point estimate of the average throughput to be within the 95% interval of confidence.

The first set of results shows the *average uplink throughput gain*, computed as the ratio between the uplink throughput achieved when CoMP is active and the uplink throughput achieved when the CoMP functionality is disabled. We selected a setting with CPP=30% and we tested the four scenarios listed in the previous section. The top plot in Figure 2 shows the uplink throughput gain for all users in the system: as expected, CoMP results in an overall decrease of system throughput, especially for high CoMP margin. To wit, if the CoMP margin is low, few users end up using CoMP, and the throughput of UEs in the center of the cell (who are unlikely to be candidates for CoMP) is unaffected, tipping the balance toward smaller throughput decrease. Recall, indeed, that CPP resources are preempted by CoMP users, but are scheduled for non-CoMP users if they go unclaimed due to the lack of potential CoMP users. Instead, for higher values of the CoMP margin, the overall throughput is smaller, especially for the Het scenario (blue
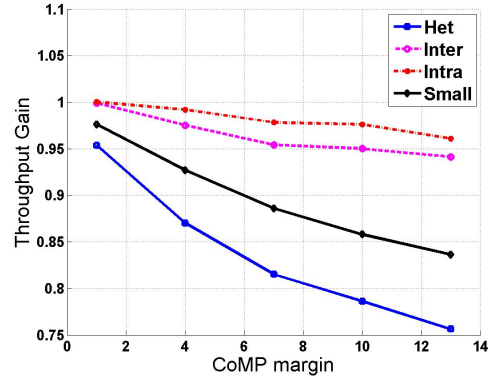
(a) All users
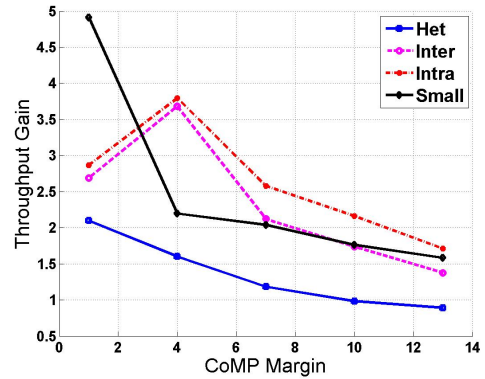
(b) CoMP users

Fig. 2. Throughput Gain for all users (a) and for CoMP users (b) as a function of CoMP margin for different scenarios (M=15, S=50, CPP=30%).

curve), where losses in excess of 20% are recorded. The reason for such a poor performance is to be sought in the non-optimized resource split among macro and small cells and in the specific user distribution in our scenario. Indeed, having many users clustered around small cells (hotspots) should call for different CPP for the two types of cells: specifically, a lower CPP for small cells (because they are serving cells for a large number of users and thus need the resources for them) than for macro cells. Restricting CoMP to small cells (black curve) does not improve the burden for small cells, but at least frees resources for macro cells, which can thus provide higher throughput when offering "umbrella" coverage to users that cannot access CoMP. The Inter and Intra scenarios, leaving small-cell resources untouched, do not feature such a dramatic throughput loss.

While these results may appear to hardly justify the use of CoMP, a different picture is painted by the bottom plot of Figure 2, showing the gain for those users only who actively use CoMP. The situation compared to the previous plot is complementary: when the CoMP margin is small, few users can use CoMP, but they have uncontested access to a large set of resources set aside for them.

In order to further investigate the remarkable gains of CoMP users in the Inter and Intra scenario, we have looked at the number of users that are involved in each scenario.

In the results discussed above, the number of UEs that were connected to a macro cell as their serving cell was 510, while 490 UEs were served by a small cell. The two values are comparable, regardless of the different coverage of the two cell types, thanks to UEs clustering around hotspots served by small cells. We also remark that these numbers of UEs are common to all four scenarios, because the choice of the serving cell is not affected by any of the CoMP policies. We thus need to look at the Het scenario to understand the number of users who select either a macro or a small cell as their CRP cell (Figure 3). As expected, the Het scenario features the largest number of CoMP users, most of which select a macro cell due to its larger coverage. The Small scenario, where only small cells can be CRP, collects approximately the same number of small cell CRPs as the Het at low CoMP margins, but the number increases with the CoMP margin that compensates for the power differential toward low-power small cells (and without any competition from macro-cell CRPs). Both Inter and Intra scenarios involve a limited number of users, which further explains the higher gains observed in Fig. 2.
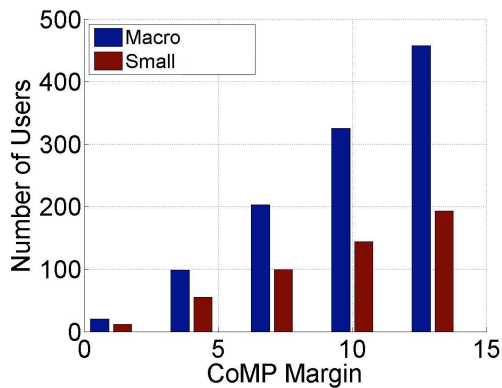
Fig. 3. Number of users connected to a specific cell type as a CRP cell, for the Het scenario.

We now turn our attention to the effects of the CoMP Pool Percentage, which are shown in Figures 4 and 5. The top two plots in Figure 4 refer to the Het scenario and detail the throughput gain as a function of the CoMP margin, for CPP values ranging from 0% to 50%. In general, the combination of high CoMP margin (which, as we have seen, increases the number of UEs eligible for CoMP) and high CPP negatively affects the throughput of both non-CoMP and CoMP users. Acceptable values of throughput loss for non-CoMP users, combined with a sizeable gain for CoMP is found for values of CPP around 20%/30% and low CoMP margin. Similar observations hold for the Small scenario in the two bottom plots of Figure 4.

We highlight an interesting behavior in Figure 5, depicting the Inter (top two plots) and Intra (bottom two plots) scenarios. For very low CoMP margin (equal to 1 dB, corresponding to less than 10 CoMP users in either scenario, as show in Figure 3), there is practically no overall throughput loss and the gain is remarkable. However,

this is due to the limited number of CoMP users, which monopolize the entire CPP resources even when CPP=10%. For margin equal to 4 dB, the number of CoMP users increases (though it is still less than 30) and starts eroding the overall throughput. The gain for CoMP users keeps rising, since CPP resources are still enough to satisfy the limited number of UEs. However, when the margin is higher than 4 dB, the overall throughput inexorably plunges, while the CoMP gain become less pronounced. Again, these trends confirm that choosing a small CPP and CoMP margin is the most balanced choice.

VI. Conclusions and Future Work

Among Coordinated MultiPoint techniques, Uplink CoMP is especially appealing since its implementation does not require UEs to be aware of it. In this paper, we have studied several configurations for Uplink CoMP in LTE-Advanced networks, using a system-oriented simulator that applies a reasonably faithful physical-layer model accounting for path loss and fast fading effects. We have introduced some parameters of interest, such as the CoMP margin and CoMP Pool Percentage, that can be used as gauges to increase the throughput of edge users. We have pointed out, however, the negative effects that incorrect settings of these parameters can determine in some scenarios. Future work will address the need to identify the optimal combination of CoMP margin and CPP.

References

[1] K. M. Karakayli, G. J. Foschini, R. A. Valenzuela, "Network Coordination for Spectrally Efficient Communications in Cellular Systems," IEEE Wireless Commun., 13(4): 56-61, Aug. 2006.
[2] J. Lee et al., "Coordinated Multipoint Transmission and Reception in LTE-A Systems," IEEE Comm. Mag., 50(11):44-50, Nov. 2012.
[3] R. Irmer et al., "Coordinated Multipoint: Concepts, Performance, and Field Trial Results," IEEE Comm. Mag., 49(2): 102-11, Feb. 2011.
[4] A. Simonsson, T. Andersson, "LTE Uplink CoMP Trial in a HetNet Deployment," IEEE VTC Fall 2012, Québec City, Canada.
[5] L. Li, P. Butovitsch, "Uplink Comp and the Applications in LTE Heterogeneous Networks: Principles and the Field Trial," IEEE ICCT 2013, Guilin, China.
[6] Y. Li, Z. Huang, "Performance of LTE-A Uplink with Joint Reception and Inter-Cell Interference Coordination," IEEE DASC 2013, Chengdu, Sichuan, China.
[7] P. Marsch, G. Fettweis, "Uplink CoMP under a Constrained Backhaul and Imperfect Channel Knowledge," IEEE Trans. on Wireless Comm., 10(6): 1730-1742, Jun. 2011.
[8] Y. Zhang, P. Hu, "A Novel Coordinated Scheduling Algorithm for Uplink CoMP Systems," IEEE ICT 2014, Lisbon, Portugal.
[9] ITU-R Recommendation M.1225, "Guidelines for evaluation of radio transmission technologies for IMT-2000," 1997.
[10] 3GPP TS 36.814, Further Advancements for E-UTRA Physical Layer Aspects.
[11] 3GPP TS 36.214, LTE E-UTRA Physical Layer.
[12] S-E. Elayoubi, O.Ben Haddada, B. Fourestie, "Performance Evaluation of Frequency Planning Schemes in OFDMA-based Networks," IEEE Trans. on Wireless Comm., Vol. 7, no. 5, May 2008.
[13] L. Suk-Bok Lee, I. Pefkianakis, A. Meyerson, X. Shugong, L. Songwu, "Proportional Fair Frequency-Domain Packet Scheduling for 3GPP LTE Uplink," IEEE INFOCOM 2009, Rio de Janeiro, Brazil.
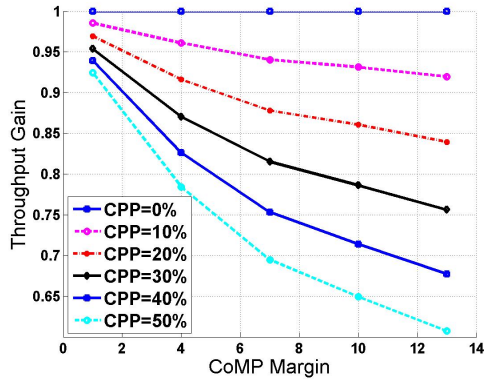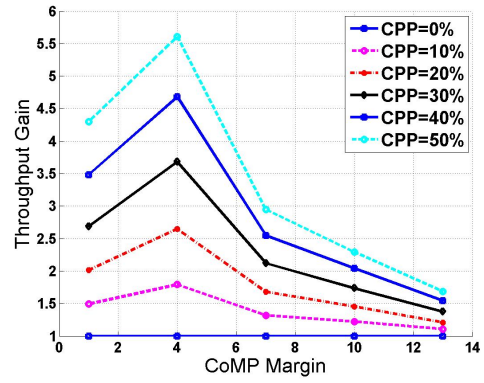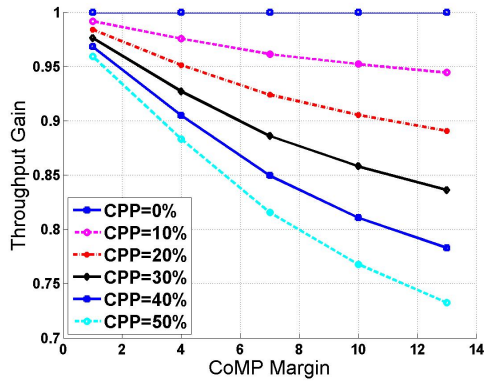
(a) All users, Het
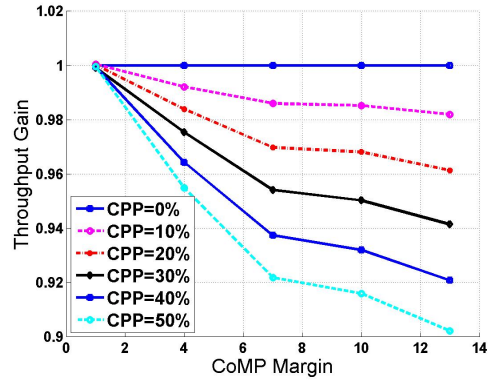
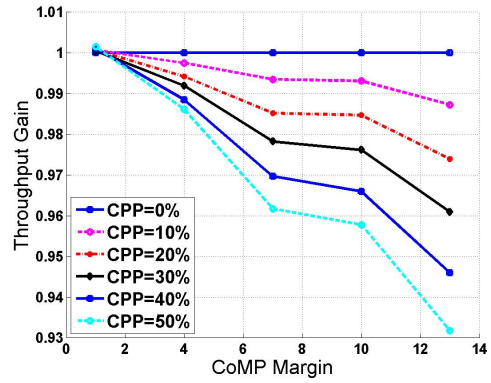
(a) All users, Inter
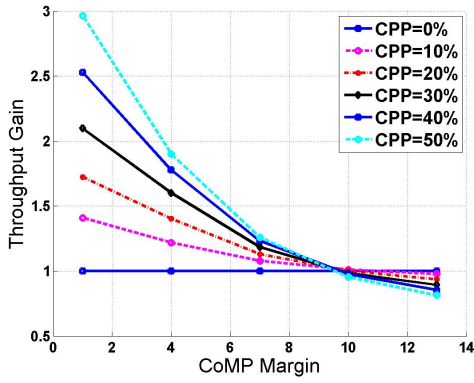
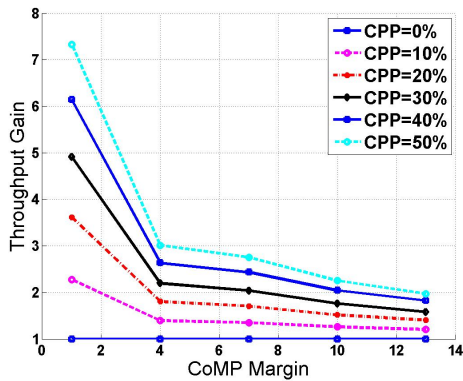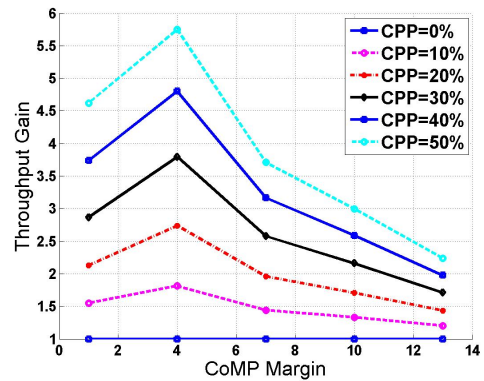(b) CoMP users, Het

(b) CoMP users, Inter

(c) All users, Small

(c) All users, Intra

(d) CoMP users, Small

(d) CoMP users, Intra

Fig. 4.   Throughput Gain as a function of CoMP margin (dB) in the Het and Small scenarios.

Fig. 5.   Throughput Gain as a function of CoMP margin (dB) in the Inter and Intra scenarios.