Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.

MODELING AND EVENT-DRIVEN SIMULATION OF COORDINATED MULTI-POINT JOINT TRANSMISSION IN LTE-ADVANCED WITH CONSTRAINED BACKHAUL

Matteo Artuso
Henrik Lehrmann Christiansen

DTU Fotonik, Department of Photonics Engineering
Technical University of Denmark
Ørsteds Plads, Building 343
Kongens Lyngby, 2800, DENMARK

ABSTRACT

Inter-cell interference (ICI) is considered as the most critical bottleneck to ubiquitous 4th generation cellular access in the mobile long term evolution (LTE). To address the problem, several solutions are under evaluation as part of LTE-Advanced (LTE-A), the most promising one being coordinated multi-point joint transmission (CoMP JT). Field tests are generally considered impractical and costly for CoMP JT, therefore the need to provide a comprehensive and high-fidelity computer model to understand the impact of different design attributes and the applicability use cases. This paper presents a novel approach to the modelling and simulation of an LTE-A system with CoMP JT by means of discrete event simulation (DES) using OPNET modeler. Simulation results are provided for a full buffer traffic model and varying the characteristics of the interface between cooperating points. Gains of up to 120% are achieved for the system throughput when using CoMP JT.

1 INTRODUCTION

The demand for always more performing mobile networks is increasing steeply. With the deployment of a next-generation global technology such as LTE, the 3rd generation partnership project (3GPP) aims at providing a new and scalable system to address these needs. In an LTE system, the radio spectrum in the downlink – and therefore the overall capacity of a cell – is shared among all the users using orthogonal frequency multiple access (OFDMA). Thanks to this, a user equipment (UE) can be scheduled in both time and frequency with a dense granularity. The smallest unit that can be assigned is represented by a physical resource block (PRB) that consists of 180 kHz in frequency and 1 ms in time (Figure 1).

The eNodeB assigns resources by means of the medium access control (MAC) scheduler. This entity can selectively assign the PRBs to the UEs depending on an advanced scheduling algorithm that typically takes into account the reported channel quality indicator (CQI) by the UE, the quality of service constraints of the traffic and a measure of fairness within the cell. The scheduling decisions are taken every millisecond and broadcasted to the UEs by means of control channels. Thanks to this, the UEs can listen only to those sub-carriers where they have been scheduled (Baker, Sesia and Toufik 2011).

3GPP decided for the adoption of OFDMA in order to provide a high degree of flexibility for different deployment scenarios, network conditions and design choices.

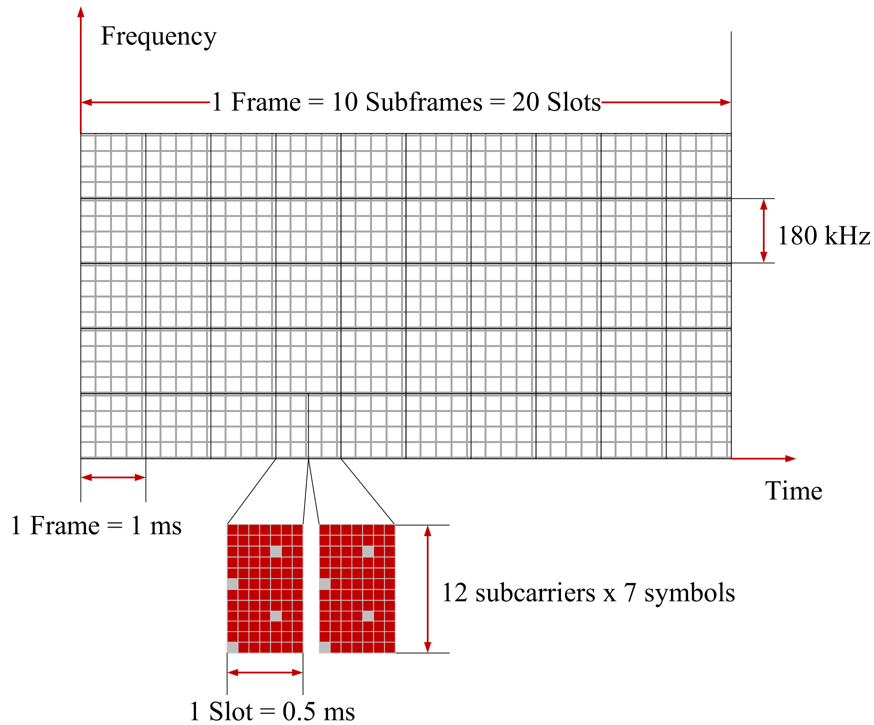

Figure 1: Resource Grid Scheme for OFDMA

One of the major bottlenecks for providing a ubiquitous and high-performance broadband access is the ICI. This is due to the reuse of the same frequency in adjacent cells, which implies high levels of noise at the cell edges, therefore limiting the achievable user throughput and the coverage of the cell itself (Baker, Sesia and Toufik 2011).

The next release of LTE – LTE Advanced – introduces a series of advanced techniques to ensure even higher performances in terms of user throughput and coverage.

In order to cope with the interference, advanced schemes for base stations cooperation are under evaluation. The most promising one – for the performances that can be theoretically achieved for cell-edge UEs – is considered to be CoMP JT (Fettweis and Marsch 2011). This technique requires the coordinated transmission of user data from multiple eNodeBs that belong to a cooperative cluster. In order for this technique to be effective, the multi-UE scheduling done at the MAC sub-layer at each eNodeB must be coordinated within the cooperating cluster. Furthermore, when a UE is attached to an eNodeB, its data from the internet is transmitted through the evolved packet core (EPC) at the S1 interface of that particular eNodeB – see Figure 2 – and therefore is not available at the other cooperating eNodeBs. The eNodeB that is serving the UE is supposed to forward the data to the other transmission points (TPs). Together with the data, the serving eNodeB must provide the identity of the UE to serve as well as the CQI value reported by the UE at the serving eNodeB. This in order for the cooperating eNodeBs to adopt the most suitable modulation and coding scheme (MCS) for the UE. This transmission is done by means of the X2 interface that exists between each pair of eNodeBs (Figure 2).

The X2 interface is not necessarily a direct physical point-to-point connection. In most implemented cases it is distributed over several links and layer 2 switches. This network between eNodeBs is typically referred to as backhaul.

We discussed above on the need of tight synchronization to exploit the benefits achievable with CoMP JT. This synchronization comprehends coordinated MAC scheduling decisions, time and frequency synchronization, up-to-date CQI information and distributed availability of user data at all the

TPs. Therefore, the bandwidth and the latency introduced by the backhaul are of primary importance for the feasibility of CoMP JT in real scenarios (Fettweis and Marsch 2011).

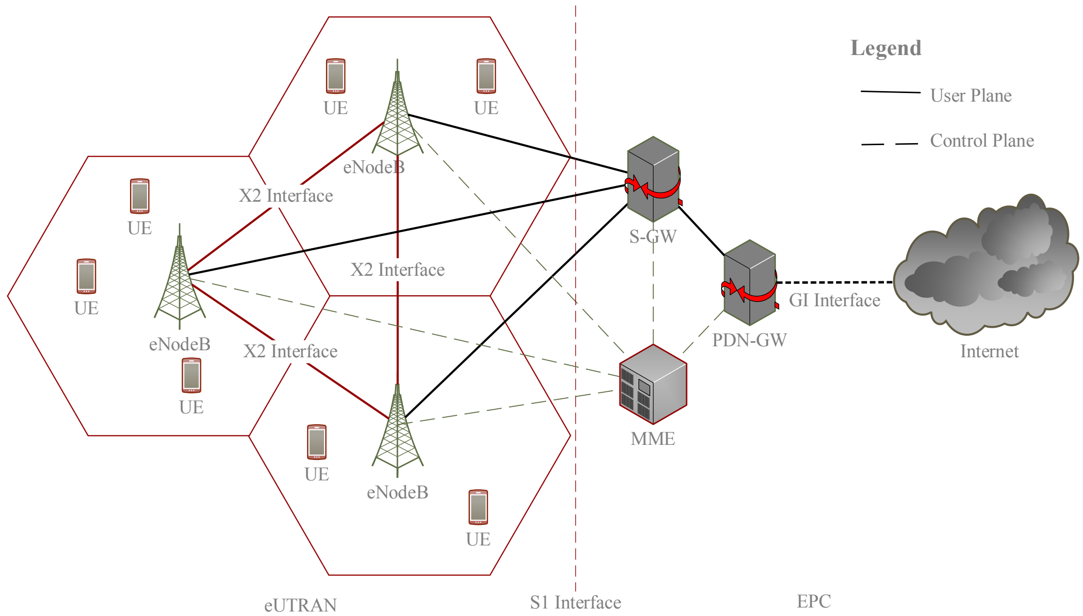

Figure 2: LTE Network Architecture

The scope of this paper is to present a novel approach to the study of CoMP JT and in particular on the impact of the characteristics of the backhaul on the gain that can be achieved. Our results are obtained by means of modelling and simulations using OPNET modeler. To ease the readability of this paper, a table with the main acronyms used is provided in Table 3 at the end of the paper.

2 RELATED WORKS

The work of Fettweis and Marsch (2011) is one of the most comprehensive studies on CoMP as it covers both theoretical and implementation aspects. The majority of the research for CoMP is in the field of modelling and simulation. Typically, the investigations carried out in these works are based on mathematical modelling of the relevant characteristics of the LTE network and on the use of standardized models for the radio transmission (Fettweis and Marsch 2011). The models are then implemented and used for simulations using – among others – Matlab and its libraries. In the following are listed some of the research works that are more related with this paper. Amin et al. (2010) present in their paper the results for an architecture very close to the one analyzed in this paper with centralized scheduling and JT and will be used as a reference for our results. Clerckx et al. (2012) provide a comprehensive analysis of the main CoMP techniques, comparing the relative gains in both homogeneous and heterogeneous deployments. Han et al. (2013) present practical design choices in real scenarios, as is done by Fettweis and Marsch (2011). The impact of the backhaul network on the performances of CoMP is investigated in terms of technology to be used by Biermann et al. (2011) and also in terms of topology and clustering techniques by Biermann et. Al (2013). The results of Molisch, Yang and Zhang (2012) present the dynamic switching between different CoMP schemes to adapt to the characteristics of the backhaul.

As regards the standardization process, 3GPP standardized the main procedures and the signaling necessary for CoMP, but left the practical implementation of the schemes open for research (3GPP 2013a). The research has been coordinated by 3GPP that provided the overall evaluation parameters and standardized scenarios to be used (3GPP 2013a). Moreover, 3GPP also standardised the framework for the investigation of the impact of the backhaul network by listing its main characteristics (3GPP 2013b).

The novelty of this paper is the use of the DES tool OPNET Modeller for modelling, implementation and simulation as an alternative to the approach of the related works. The modelling will be carried out using the characteristics of the framework standardised by 3GPP. The impact of the backhaul network will be analyzed in terms of the latency introduced by this network as this represents the major issue for CoMP JT as pointed out in the aforementioned related works and others.

3 SYSTEM MODEL

3.1 System Architecture

The modelled system refers to the scenario 2 in (3GPP 2013a) where 3GPP depicts a homogeneous network with a central eNodeB and a surrounding cluster of high-power remote radio heads (RRHs) connected via optical fiber. The main difference in our model consists in the fact that all the TPs are macro eNodeBs, therefore modelling a more distributed network that is likely to be found in urban areas with inter-site distance (ISD) of 500 m and transmitting power of 46 dBm.

As can be seen in Figure 3, the model adopts the typical hexagonal shape for the cells with one eNodeB per cell. The cluster is static with 3 eNodeBs with omnidirectional antennas. The area with the highest level of ICI is the center of the cluster.

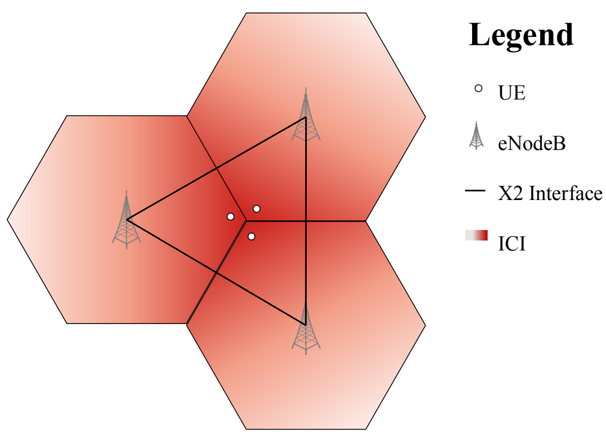

Figure 3: System Architecture Model

The model is focused on the radio access part of an LTE network, therefore the EPC is simplified and the S1 interface points directly towards an external traffic generator that models a full-buffer traffic model. The X2 interface is modelled by means of a variable attribute that represents the latency experienced by the transmissions between the TPs in the cluster.

There are two kinds of nodes in the network, eNodeBs and UEs. Both node models have an internal structure that resembles the lower layers of the LTE protocol stack (Baker, Sesia and Toufik 2011) with physical and MAC layers for the UE, physical, MAC and radio link control (RLC) layers for the eNodeB. The eNodeB moreover keeps the physical and the MAC layers separated depending on the network interface they refer to – X2, S1 or radio. The modular structure for the nodes allows extendibility and customizability of the model in terms of additional features without affecting the overall design. Each module is designed as a finite state machine (FSM) to ease the implementation of protocols. Even though most of the simulated scenarios have static UEs, the UE can be mobile, but its trajectory does not take it outside its initial cell since handover is not implemented.

3.2 eNodeB Model

The packets from the traffic generator are received at the S1 interface and forwarded to the RLC sub-layer that segments them according to the instructions provided by the MAC sub-layer (Figure 4). These instructions consist on the value of bits per OFDM symbol that can be used for the PRBs assigned to that particular UE. This value is based on the latest available CQI value reported by the UE at the serving eNodeB. The number of OFDM symbols per PRB is fixed to 168 in our model, since 8 symbols are used for the reference signals (RSs) as per Figure 1.

The RLC sub-layer passes the segments to the MAC sub-layer for enqueueing and scheduling. In case CoMP JT is enabled, then the segments are also copied to the X2 interface. The modules for the layers at the X2 interface forward a copy of the segments to each collaborating eNodeB including CQI and ID of the UE to be served. Self-organizing network (SON) techniques are implemented in the modules of the X2 interface in order to discover the topology of the cluster and the identity of the collaborating eNodeBs.

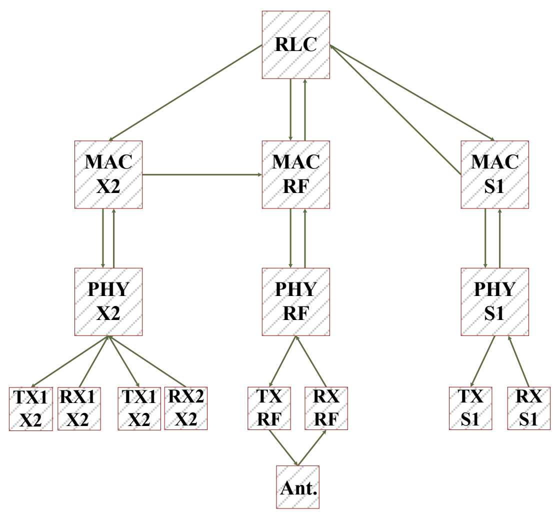

Figure 4: eNodeB Node Model

On the radio interface of the eNodeB, the scheduler is implemented at the MAC sub-layer. Every scheduling turn, the algorithm allocates PRBs to the UEs in a round robin manner, but also evaluating the size of the transmission sub-queue and the CQI value for that UE. The study is limited to the use of one scheduling algorithm to remove this degree of variation from the comparison between the cases with and without CoMP. Once the scheduling is completed, the MAC model associates the IDs of the UEs scheduled in correspondence with the PRB index to model the physical download control channel (PDCCH). This list of scheduling assignments is then accessed remotely by the UE to simplify the modelling of the PDCCH. In case CoMP is enabled, the scheduling decisions are coordinated within the cluster. This implies that a UE is scheduled in the same PRB by all the cooperating TPs in the cluster. This is necessary to successfully exploit the interference cancellation as signal quality enhancement. The coordination of the schedulers is modelled in two steps: the eNodeBs discover the topology within their own cell and afterwards they access the list of UEs of the neighboring cells to have a complete list of the UEs in the cluster. This overall list is then used by each scheduler to allocate the PRBs.

All the scheduled segments are passed to the physical layer. Each segment contains the index of the PRB where it is going to be transmitted and the MCS index decided by the MAC, so the physical layer can set the correct modulation on the correct sub-carriers. The main characteristics of the sub-carriers are set statically at the initialization of the simulation by the module of the physical layer. Among these: the bandwidth (15 kHz), the inter-carrier spacing (~ 1.7 kHz), a reference value for the data-rate of the channel (1 Mbps) and the minimum frequency for each channel –we modelled 20 MHz bandwidth in LTE band 7, therefore from 2620 MHz to 2640 MHz (Baker, Sesia and Toufik 2011).

A further relevant static setting regards the receiving channels that each transmitting channel shall evaluate. In fact, in order to model the broadcast nature of radio transmissions, whenever a segment is transmitted on a channel, the segment is copied to all the receving candidates listed in the transmitting channels. This list can be built dynamically each time, evaluating all the possible transreceiver pairs in the network. This evaluation can severely affect the performances of the simulation due to the large number of channels in OFDM systems. Therefore, in order to optimize the model and design it closer to reality, each transmitting channel has a static list only of the corresponding receiving channels in the UEs in the entire network at the same frequency. A segment meant for a specific PRB is further segmented to follow the structure of the 12 separated channels and RSs are inserted to model the pattern of Figure 1. The RS are modulated with quadrature phase-shift keying (QPSK) for robustness, therefore each OFDM RS symbol can carry 2 bits (Baker, Sesia and Toufik 2011). They are modelled with a dedicated signal of 16 bits that is segmented into 8 parts that the UE then reassembles.

The modelling of the radio transmission as well as the procedures of the physical layer are partially carried out by the module within the eNodeB model and partially by the radio transmission model.

3.3 Radio Transmission Model

The radio transmission is modelled with a sequence of functions that evaluate each segment once it has been transmitted by the eNodeB (Figure 5). Each of the functions models a different aspect of the radio channel, taking into account the parameters of the nodes involved in the transmission as well as those that are interfering with it. The results of these functions are used as labels on the segment to be further used by the following functions in the sequence. The first functions of the sequence set the transmission delay, the propagation delay and the parameters related to the closure of the link and the matching of the channels, which are always successful as they refer to the static list built by the physical layer within the eNodeB. The antenna gains in both transmitter and receiver are set to 0 dB since we modelled omnidirectional antennas.

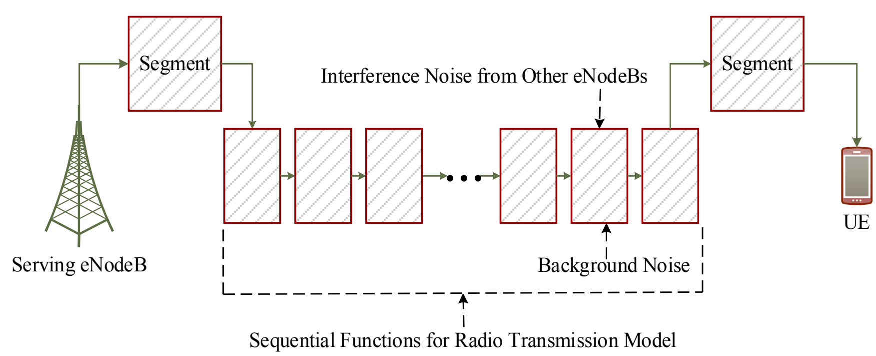

Figure 5: Radio Transmission Model

The most relevant functions are those that perform the evaluation of the nature of the incoming segments and the interference-management operations. Whenever a segment arrives from the serving eNodeB in one of the receiving channels of the UE, then a function checks whether the current channel belongs to a PRB scheduled to the UE. If that is the case, then the segment continues the processing without any further modification, the channel is locked in order to receive properly the current segment and the received power is calculated. In case the channel does not belong to a PRB assigned to the UE, then the segment is labelled as noise since it is not meant for that UE, but for another one in the cell. When the segment received has not been transmitted from the serving eNodeB, then it is necessary to distinguish the following cases: in case CoMP JT is not enabled, then the segment is labelled as noise and its received power will contribute to increase the overall noise of the other transmissions. In case CoMP JT is enabled, then the scheduling assignments of the transmitting eNodeB are checked, similarly to the way it is done for the serving eNodeB. In case the UE is scheduled in that particular PRB from the cooperating eNodeB, then the segment is labelled as noise, but with CoMP gain. In fact, these particular noise segments in reality will enhance the SINR of the segment transmitted from the serving eNodeB since they are a copy of it. Finally, in case the UE is not scheduled in the particular PRB to which the receiving channel belongs to, than the segment is labelled as noise without CoMP gain.

The ICI is modelled with these differences in the way that incoming segments are labelled. Whenever a segment is received on a channel where another segment is already being currently received, their received power is used differently in the calculation of the SINR, depending on these labels. In case the reception of a valid segment – i.e. from the serving eNodeB – is affected by an incoming noise segment, then the received power of the noise segment contributes to the interference noise experienced by the valid segment. Nevertheless, if CoMP JT is enabled and the segment is labelled as CoMP gain, then the noise segment is in fact a copy of the valid segment transmitted from a cooperating eNodeB in the cluster. Thanks to this, the received power of the noise packet contributes to increase the received power of the valid packet, rather than contributing to the interference noise. A further degree of freedom is represented by the capabilities of the UE to convert successfully the interfering power of a cooperating segment. In fact, different UE types could be capable of converting only a fraction of the overall power of the cooperating segment. This is modelled by means of a parameter that acts as a weight to balance the fraction of power that is successfully converted and the one that is not – therefore contributing to the interference noise. An evaluation of the SINR in watt is presented in (1), where P_{vs} is the received power of the valid segment, P_{cs} the received power of the cooperating segment, $\alpha \in [0,1]$ is the parameter to model the capabilities, N_b is the background noise at 290 K and N_i the interference noise already experienced by the valid segment:

$$SINR_W = \frac{P_{vs} + \alpha P_{cs}}{N_b + N_i + (1-\alpha) P_{cs}}. \quad (1)$$

It is possible to notice that when $\alpha = 1$ then the UE converts perfectly the entire received power of the cooperating segment to useful signal and there is no contribution to N_i. This is the value used for our simulations with CoMP JT enabled. On the other hand, when $\alpha = 0$, then the UE is not capable to convert the interference and this corresponds to the case when CoMP JT is not enabled. Once the SINR evaluation for a valid segment is completed, the noise packets are no longer processed. The SINR is recalculated every time a valid reception is interrupted by a noise segment and the values of the received power and the interference noise are updated at every turn. The result of (1) in dB is used in (2) where g_p is the processing gain of the modulation set at the UE for that particular channel to calculate the effective SINR:

$$SINR_{dB}^{eff} = SINR_{dB} + g_p. \quad (2)$$

The value of the effective SINR is then used to have the correspondent value of the BER (Bit Error Rate). In order to do that, OPNET provides a set of functions that take as input the SINR calculated by the UE and the modulation used for the transmission to calculate the correspondent BER. The calculation uses the modulation tables and curves correspondent to the one used as input. The value of the BER is then used to calculate the number of bits in the segment that shall be considered as error bits due to incorrect reception. We used a uniform distribution in [0,1] together with the BER value to compute the number of errors. It is possible to drop segments that exceed the capabilities of error correction of a receiver. These capabilities are modelled with an attribute that varies in [0,1] where 0 means that only segments with 0% of error bits should be accepted at the receiver and 1 that even segments with 100% error bits can be accepted. We statically assigned this attribute to 1, since we are interested in receiving all the packets within the node model of the UE for further processing. Once this last function is completed, the channel for that particular sub-carrier is unlocked and the sequence of functions is repeated for the next segment in the transmission queue.

A particular case is represented by signals that are broadcasted from the eNodeBs in order to provide synchronization information and cell identities. These functionalities are modelled with a dedicated signal sent periodically by the eNodeB to all the UEs on the central frequency (sub-carrier with index 599). These segments are treated differently because of their importance and priority. When a synchronization segment is received, it is processed further only if it has been transmitted by the serving eNodeB, otherwise it is ignored. As said, the synchronization signal is used also to carry the identity of the serving eNodeB to the UE in order to make possible the correct detection of the following segments. Therefore at the simulation start, the first synchronization signal sent, is treated differently from the following since the UE does not have the identity of its serving eNodeB yet. In order to detect the correct one to attach to, the UE evaluates its relative position within the cell and the cluster in relation to the eNodeBs that transmit these initial signals. The eNodeB geographically closer to the UE at the simulation start, is picked as the serving one.

3.4 UE Model

The segments are retrieved by the physical layer of the UE after the completion of the sequence of functions that models the radio transmission. This module records the statistics that are used for the comparison: for each data segment, it records the number of bits successfully received and then stores the segment in a reassembly buffer. In case the segment is part of a RS, the module records the SINR experienced by that segment. The UE shall also refresh the list of PRBs where it is scheduled and this is done right after the completion of the scheduling operations at the eNodeBs. Concurrently, the reassembly buffers are emptied for the next turn and the complete packets move to the MAC layer. The SINR and BER experienced by the complete packets of the RSs (see above) are used to perform the link adaptation. The module determines the current CQI index and the most suitable modulations in order to have a BLER (Blocking Error Rate) of at most 10%. According to the CQI reporting period set by the eNodeB, the MAC calculates multiple values of the CQI and then it feeds back their average. The MAC layer receives also complete data packets to provide compatibility with future extensions of the model.

4 SIMULATION RESULTS

In all the simulated scenarios, we used the throughput measured at the physical layer of the UE. This metric considers only the number of successful bits received. This performance metric is used to compute both the throughput perceived by a cell-edge UE and the cell/system throughput. The CoMP gain in percentage is calculated as per (3) where T_c is the throughput with CoMP JT and T_{nc} the one without:

$$g_{CoMP}^{\%} = \frac{T_c - T_{nc}}{T_{nc}} * 100. \tag{3}$$

4.1 Impact of Cell Load and UE Distributions

Simulation parameters for this set of simulations are provided in Table 1 and the results, together with the distribution of the UEs are presented in Figure 6. The UEs are distributed in two different patterns in case A (blue) and B (red). The system throughput is calculated as the summation of the throughput experienced by each UE.

Table 1: Simulation Parameters for First Set

Simulated Time	5 s	Synchronization Period	20 ms
X2 Latency	0 s	UEs per cell	1, 3, 5, 10 in 2 distributions (A and B)
CQI Reporting Period	10 ms	Performance Metrics	System Throughput & Throughput Gain
ISD	500m		

As can be seen in Figure 6, CoMP JT is beneficial when the distribution of the UEs is towards the center of the cluster where the ICI is at its highest, as per Figure 3. CoMP JT introduces a considerable gain for the system throughput in the case A of ~50% for up to 5 UEs per cell and of ~80% in the scenario with 10 UEs per cell. Using CoMP in areas with good radio quality does not generally introduce significant gains, but it can even be a loss on a system level as can be seen in case B for 5 and 10 UEs. The design of system-level CoMP JT should therefore be selective and based on the reporting of the UE. The serving eNodeB shall therefore trigger the cooperation procedures within the cluster when the CQI from a certain UE – or a pool of UEs – goes below a certain threshold. This in order to keep the gains that can be obtained on a system level, without encountering the loss of non-beneficial use of CoMP JT.

A reasonable solution might therefore be a continuous exchange of messages related with the quality of the signal within the cluster, with the UE reporting also measurements from neighboring eNodeBs. The serving eNodeB is therefore able to change dynamically the topology of the cluster communicating with the eNodeBs that can be beneficial for the JT to that UE. Once the cluster topology is specified for that particular UE, the serving eNodeB can start the exchange of UE data and scheduling information to the other TPs as said above.

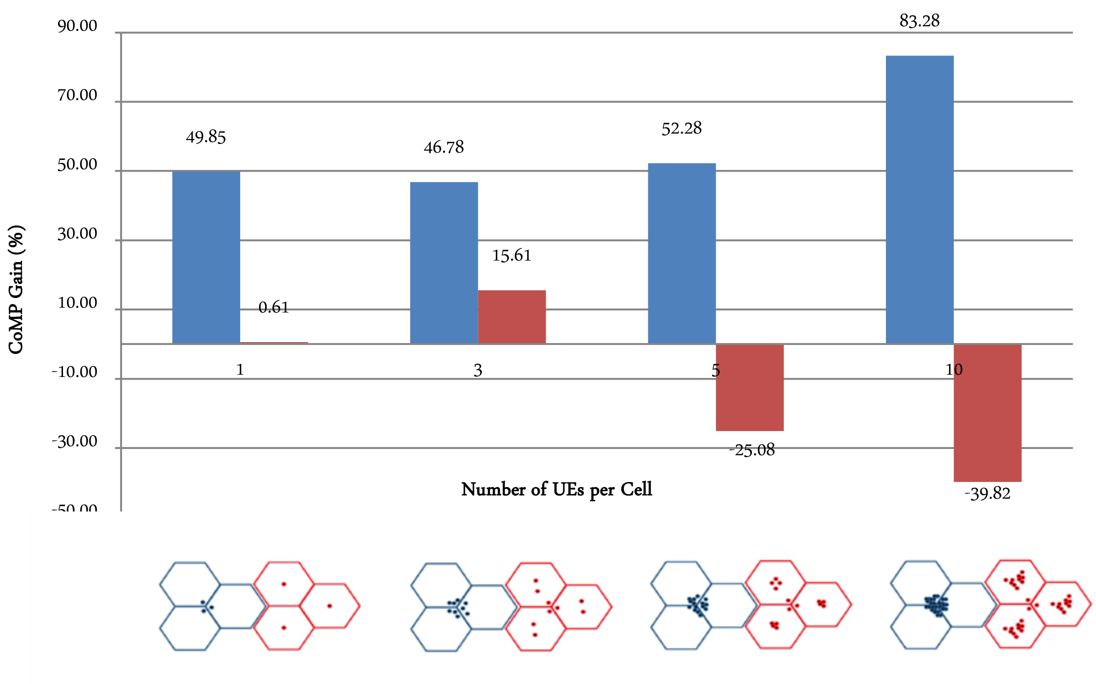

Figure 6: System Throughput Gain vs Cell Load

4.2 Impact of X2 Latency

The second set of simulations adopts the same performance metrics as above, but the scenarios have only 1 UE per cell. The UE in one cell moves from the center of its cell towards the center of the cluster at pedestrian speed (3 km/h). The simulated scenario is repeated with different values of the X2 latency. This shows how the characteristics of the backhaul can affect the feasibility of CoMP JT in cases that are closer to possible deployments than those with perfect backhaul as in the previous simulation set. Simulation parameters are provided in Table 2.

Table 2: Simulation Parameters for Second Set

Simulated Time	3 min	Synchronization Period	20 ms
X2 Latency	0 – 0.02 s	UEs per cell	1
CQI Reporting Period	10 ms	Performance Metrics	System Throughput & Throughput Gain
ISD	500 m	Speed Pedestrian UE	3 km/h

Figure 7 presents the gain of the system throughput. It is possible to notice a significant change from ca. 120% to ca. 40% when the latency is equal to the scheduling period of 1 ms. This confirms the results presented in the related works, in particular Clerckx et al. (2012) for the gain of a cell-edge UE with low latency on the X2 and Amin et al. (2010) for the impact of the latency on the system gain. Regarding this second point, Figure 8 presents the results from the same simulation set comparing the loss in terms of system gain compared to the case with an ideal interface with no latency. It is possible to notice a further clarification of what pointed out before, with the gain becoming a loss on a system level with a latency of 10 ms and more with a negative gain in Figure 7 and an impact in Figure 8 that exceeds -100%.

The results obtained in this scenario with a pedestrian UE define two main thresholds for the latency of the X2 interface that are 1 ms and 10 ms. Our future works will extend these findings by simulating UEs moving at different speeds and analyzing how these thresholds change accordingly. We expect to be able to define what is the maximum speed a UE can have to experience a gain from CoMP JT for a certain value of the X2 latency in the backhaul network.

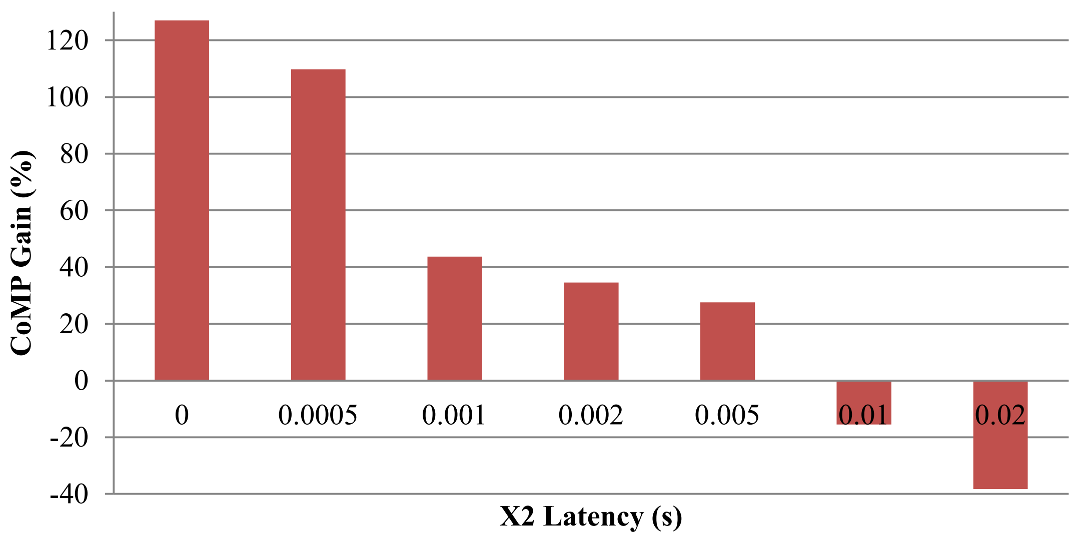

Figure 7: System Throughput Gain vs X2 Latency

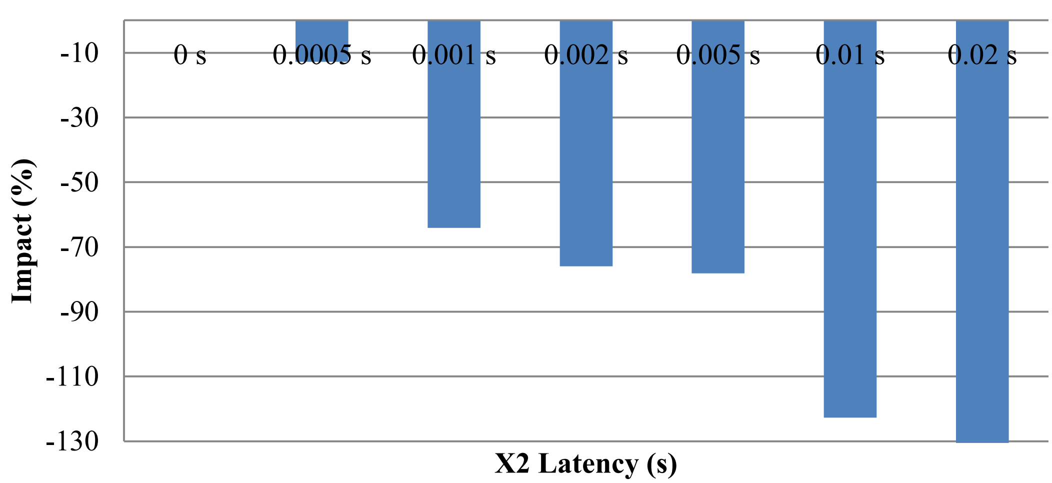

Figure 8: Impact on System Throughput Gain vs X2 Latency

5 CONCLUSIONS

In this paper we designed a modular and extendible model to be used with the DES tool OPNET modeler. It allows a thorough investigation of the impact of the latency of the backhaul network, of the cell load and of the distribution of the UEs on the feasibility of the deployment of CoMP techniques. Our results show that with CoMP JT the system throughput experiences a gain of up to 120% in the scenario with a pedestrian UE with X2 latency below 1 ms. Values of the latency below 10 ms still ensure a considerable gain from 40% to 25% proportionally. In case the latency exceeds this threshold, CoMP should not be used since it implies a loss for the network performances.

Our future works envisage the extension to advanced multiple antenna schemes, the impact of CoMP in handover scenarios, how CoMP affects the behavior of different kinds of traffic and applications and finally the modelling of a more centralized strategy for CoMP (Cloud-RAN) and its comparison with the distributed approach presented in this paper, taking into account the requirements of the backhaul network.

ACKNOWLEDGMENTS

This work was partially sponsored by the 7th Framework Programme for Research of the European Commission HARP project, under grant number HARP-318489.

A LIST OF ACRONYMS

Table A: Definition of the Main Acronyms Used in the Paper

3GPP	3rd Generation Partnership Project	MCS	Modulation and Coding Scheme
BER	Bit-Error Rate	OFDMA	Orthogonal Frequency Division Multiple Access
CoMP	Coordinated Multi-Point	PRB	Physical Resource Block
CQI	Channel Quality Indicator	PDCCH	Physical Download Control Channel
DES	Discrete-Event Simulation	QPSK	Quadrature Phase Shift Keying
EPC	Evolved Packet Core	RLC	Radio Link Control
FSM	Finite-State Machine	RRH	Remote Radiohead
ICI	Inter-Cell Interference	RS	Reference Signal
ISD	Inter-Site Distance	SINR	Signal to Interference and Noise Ratio
JT	Joint Transmission	SON	Self-Organizing Network
LTE	Long Term Evolution	TP	Transmission Point
MAC	Medium Access Control	UE	User Equipment

REFERENCES

3GPP. 2013a. "36.819 Coordinated multi-point operation for LTE physical layer aspects version 11.2.0 Release 11". Sophia Antipolis Cedex: ETSI.

3GPP. 2013b. "36.874 Coordinated multi-point operation for LTE with non-ideal backhaul version 12.0.0 Release 12". Sophia Antipolis Cedex: ETSI.

Amin, M. A., S. Brueck, J. Giese and L. Zhao. 2010. "Centralized Scheduling for Joint Transmission Coordinated Multi-point in LTE-Advanced." In *International ITG Workshop on Smart Antennas (WSA)*, 177–184. Bremen: Institute of Electrical and Electronics Engineers, Inc.

Baker, M., S. Sesia and I. Toufik. 2011. *LTE - The UMTS Long Term Evolution From Theory to Practice*, 2nd ed. Chichester: John Wiley & Sons Ltd.

Biermann, T., C. Choi, H. Karl, W. Kellerer and L. Scalia. 2013. "How Backhaul Networks Influence the Feasibility of Coordinated Multipoint in Cellular Networks." In *IEEE Communications Magazine* 51:168–176.

Biernmann, T., C. Choi, S. Mizuta and L. Salia. 2011. "Coordinated Multipoint Multiuser-MIMO Transmissions over Backhaul-Constrained Mobile Access Networks." In *IEEE 22nd International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC)*, 1336-1340. Toronto: Institute of Electrical and Electronics Engineers, Inc.

Clerckx, B., E. Hardouin, D. Lee, D. Mazzarese, S. Nagata, K. Sayana and H. Seo. 2012. "Coordinated Multipoint Transmission and Reception in LTE-Advanced: Deployment Scenarios and Operational Challenges." In *IEEE Communications Magazine* 50:148–155.

Fettweis, G. P. and P. Marsch. 2011. *Coordinated Multi-Point in Mobile Communications* Cambridge: Cambridge University Press.

Han, S., X. Hou, A. F. Molisch and C. Yang. 2013. "How Do We Design CoMP to Achieve its Promised Potential?" In *IEEE Wireless Communications* 20:67–74.

Molisch, A.F., C. Yang and Q. Zhang. 2012. "Cooperative Downlink Transmission Mode Selection under Limited-Capacity Backhaul." In *IEEE Wireless Communications and Networking Conference (WCNC)*. 1082-1087. Shanghai: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

MATTEO ARTUSO is a research assistant at the Technical University of Denmark where he is also the chairman of the IEEE student branch. He worked in the field of computer simulations during his M.Sc. and B.Sc. with projects related to next-generation mobile networks. This paper is part of his thesis project with the title "Backhaul Impact on Coordinated Multipoint in LTE Networks". His address is matart@fotonik.dtu.dk.

HENRIK CHRISTIANSEN is an Associate Professor in mobile communication at the Technical University of Denmark, where he also received his M.Sc.E.E. and Ph.D., specializing in telecommunications. He also has several years of experience from the telecom industry. His main areas of research are mobile network architectures, mobile fronthaul and backhaul networks. His address is hlch@fotonik.dtu.dk.