

Assessment of Distributed Multi-User MIMO Transmission in 5G Networks

Braam, Sjors; Litjens, Remco; Smulders, Peter; Ijntema, Wieger

DOI

[10.1145/3416012.3424629](https://doi.org/10.1145/3416012.3424629)

Publication date

2020

Document Version

Final published version

Published in

MobiWac 2020

Citation (APA)

Braam, S., Litjens, R., Smulders, P., & Ijntema, W. (2020). Assessment of Distributed Multi-User MIMO Transmission in 5G Networks. In *MobiWac 2020 : Proceedings of the 18th ACM Symposium on Mobility Management and Wireless Access* (pp. 125-132). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3416012.3424629>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Assessment of Distributed Multi-User MIMO Transmission in 5G Networks

Sjors Braam
Dept. of Networks, TNO
The Hague, The
Netherlands
sjors.braam@tno.nl

Remco Litjens
Dept. of Networks, TNO
The Hague, The
Netherlands
Dept. of EEMCS, TU Delft
Delft, The Netherlands
remco.litjens@tno.nl

Peter Smulders
Department of Electrical
Engineering, TU Eindhoven
Eindhoven, The Netherlands
p.f.m.smulders@tue.nl

Wieger IJntema[†]
Department of TT-A12,
IAV GmbH
Gifhorn, Germany
wieger.ijntema@iav.de

ABSTRACT

We present a simulation-based assessment of the performance potential of distributed MIMO (D-MIMO), multi-user MIMO (MU-MIMO) and particularly the combined D/MU-MIMO operation, for which we extend previously published scheduling and beamforming principles. The assessment study reveals that, when optimizing average user throughput performance, D-MIMO, while fruitless when used in isolation, is very effective when intelligently combined with MU-MIMO. Alternatively, when optimizing the cell edge performance, MU-MIMO, while also shown to be ineffective when used in isolation, is in fact very valuable when accompanied by a suitable configuration of D-MIMO. As an illustrative example, when using a jointly optimised configuration of D/MU-MIMO in a highly loaded urban deployment scenario, a 121% cell edge performance gain can be attained over a scenario using only D-MIMO, and even a demonstrated 153% gain over a scenario where only the MU-MIMO feature is available.

CCS CONCEPTS

• Networks • Wireless access networks • Mobile networks
• Network management • Network algorithms • Network performance evaluation

KEYWORDS

Mobile networks, 5G, MIMO, distributed MIMO, multi-user MIMO, beamforming, scheduling, performance assessment.

ACM Reference format:

Sjors Braam, Remco Litjens, Peter Smulders and Wieger IJntema. 2020. Assessment of Distributed Multi-User MIMO Transmission in 5G Networks. In *Proceedings of the 18th ACM International Symposium on*

[†]The research was done while Wieger IJntema was with TNO, The Netherlands.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

MobiWac '20, November 16, 2020, Alicante, Spain.

© 2020 Association of Computing Machinery.

ACM ISBN 978-1-4503-8119-2/20/11...\$15.00.

DOI: <https://doi.org/10.1145/3416012.3424629>

Mobility Management and Wireless Access (MobiWac '20). ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3416012.3424629>.

1 Introduction

The development of 5G mobile network technology [1][2] is primarily driven by the challenging performance requirements imposed by use cases and applications in the industrial and consumer domains, and by the demand for a higher areal capacity in support of the ever-growing mobile traffic volumes. The network capacity predominantly depends on three factors, viz. the cell density, the available spectrum and the spectral efficiency. *Cellular densification* is a costly yet very important on-going task of mobile network operators, driven also by the envisioned use of ever higher spectrum and the associated challenges in terms of more severe signal attenuation. Although challenging, the use of *high spectrum* is essentially mandated by the need for more spectrum and the reality that significant chunks of available spectrum can only be found in the higher regions. Spectrum in the 3.5 GHz and 26 GHz bands has been designated as most important for globally harmonised use in 5G networks. Lastly, the *spectral efficiency* is primarily influenced by the degree of frequency reuse, the signal strength of UEs (User Equipments) towards their serving BSs (Base Stations) and the experienced interference levels. In this light it is important to note that mobile network operators generally care less about the cell-average spectral efficiency, but rather focus their planning and optimisation efforts on the *cell edge performance*.

In 5G radio access technology, a number of advanced features can be exploited to enhance the (cell edge) spectral efficiency by enhancing one or more of the mentioned key aspects. Massive MIMO-based (Multiple Input Multiple Output) multi-antenna transmission is a key transmission technology used in 5G networks that can be exploited to address these aspects. With massive MIMO, a BS can apply *beamforming* to serve a given UE with a narrow beam with a high effective antenna gain and, consequently, a high received signal strength. This is essential for transmission at higher frequencies to overcome the correspondingly poor signal attenuation. Fortunately, the potential of beamforming is increased by the higher number of antenna elements that can be incorporated at higher frequencies within a reasonable form factor, given the typical half-wavelength antenna element spacing. Such beamforming is noted to also significantly reduce the interference footprint of a

given transmission and hence decreases the chances of another active UE to experience interference from this transmission.

Within a given cell, multi-antenna technology also enables *multi-layer transmission*, where multiple independent layers are separately beamformed (digitally precoded) if the associated channels are deemed sufficiently uncorrelated to avoid excessive inter-layer interference. Multiple layers can be transmitted to a single UE (SU-MIMO/SM ~ Single-User MIMO / Spatial Multiplexing) or to multiple UEs (MU-MIMO ~ Multi-User MIMO; see Figure 1), with both the potential to (in)directly enhance both UE throughput and cell capacity. With 5G networks generally planned with contiguous frequency reuse (reuse of 1), the application of multi-layer transmission such as MU-MIMO can be regarded as a further densification of frequency reuse to reuse factors below 1.

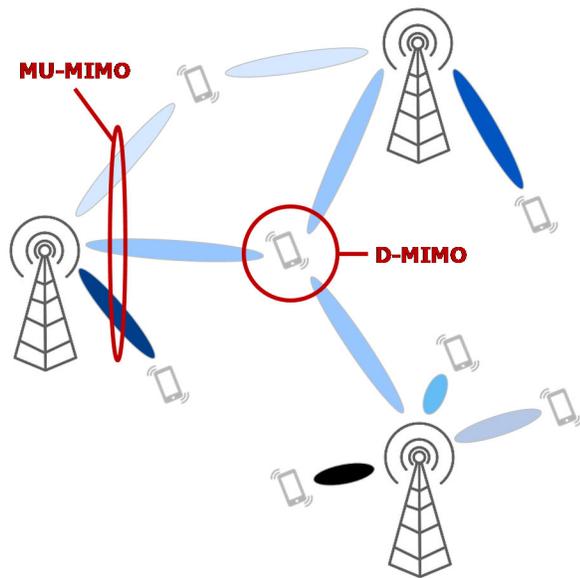


Figure 1: Concepts of beamforming and multi-antenna transmission towards multiple users (MU-MIMO) or from multiple BSs to an individual user (D-MIMO).

Also visualised in Figure 1, the third feature of relevance in this context is *distributed MIMO* (D-MIMO), which entails the collaborative transmission of multiple BSs to a given UE in the same time-frequency resources, such that the signals coherently overlap at the receiver antennas. In essence, this is no different from the beamforming as mentioned before, except that this transmission scheme jointly utilises transmit antennas available at multiple BSs, i.e. at different physical locations. In literature, this feature is known under different names, including distributed/network/virtual MIMO [1][3] in the academic literature, CoMP/JT (Coordinated Multi-Point / Joint Transmission) in 3GPP terminology [4] or cell-free networking [5] and can even be regarded as an advanced form of 3G's soft handover feature. As one can logically expect the benefits of D-MIMO transmission to be primarily experienced by UEs at/near cell edge locations, this is indeed a good candidate feature in light of the operators focus on cell edge performance or, in other words, to achieve a spatially more homogenous performance.

CoMP/JT has been the subject of many different studies. Some simulations show very optimistic results [6], but it is also known that only a fraction of the predicted gains can be reached if the CSI (Channel State Information) feedback is modelled realistically [7]. CoMP with limited and/or quantised feedback has been studied in [8]. The shift towards TDD (Time Division Duplexing) operation of mobile networks is particularly attractive for CoMP/JT and other techniques that heavily depend on CSI, like MIMO and beamforming. In TDD, the same spectrum is used for the downlink and uplink and therefore reciprocity between the downlink and uplink channels can be exploited. The CSI can be determined from uplink pilot signals, which removes the CSI feedback delay and the necessity for CSI quantisation. In [9], CoMP/JT is evaluated in a TDD system, where performance gains of more than 20% are observed for cell edge users. However, infinite-capacity and zero latency links between the BSs are assumed.

The objective of this paper is to do a coherent simulation-based assessment of the potential of SU/MU-MIMO and D-MIMO on the downlink throughput performance, where we address these technological concepts in isolation for reference purposes but primarily focus on their combined exploitation. The assessment will concentrate on both average and cell edge performance effects. We will derive the sensitivity of the optimal configuration of the MU-MIMO and D-MIMO features with respect to distinct scenario aspects, thereby shedding light on the robustness of these configurations. Elaborate discussions will be included to qualitatively understand all key observations.

The outline of the remainder of the paper is as follows. A number of key modelling aspects will be described in Section 2, while all aspects related to radio resource management are described in separate Section 3. Section 4 then outlines the assessment scenarios and defines the addressed KPIs (Key Performance Indicators). Subsequently, the numerical results will be presented and thoroughly analysed in Section 5. Section 6 ends this paper with a statement of the key conclusions and recommendations for further research.

2 Model aspects

This section describes the key modelling aspects relevant for the assessment study, considering network aspects, the user/traffic characteristics and the propagation environment.

We consider a *network* comprising 19 3-sectorised sites deployed in a hexagonal layout, as depicted in Figure 2, with an environment-specific inter-site distance for the considered urban and rural environments (see Table 1). Although all $19 \times 3 = 57$ cells are simulated in equal detail, performance statistics are collected only for the indicated 7×3 'evaluation cells'. All other cells are included in order to avoid boundary effects, to establish a realistic interference environment and to ensure the availability of antennas for D-MIMO for users located near the edges of the evaluation area. Denote with $\mathcal{C} = \{1, \dots, 57\}$ the set of all cells.

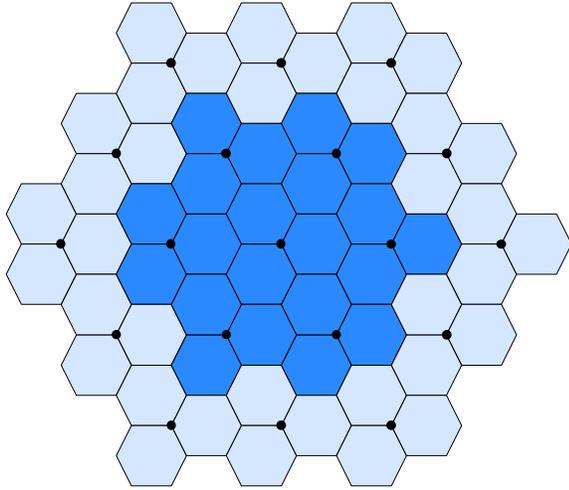


Figure 2: The assumed hexagonal network layout comprises 19 3-sectorised sites. Performance statistics are collected for the 7 × 3 more darkly shaded central cells.

All cells are uniformly equipped with either a 16T16R or a 64T64R antenna array (denoting $N_T = 16, 64$, respectively), all comprising $8 \times 8 \times 2 = 128$ cross-polarised antenna elements (AEs) with half-wavelength inter-AE spacing, but with different vertical subarray sizes. The individual AEs are modelled as specified in Table 7.3-1 in [10] with a maximum per-AE gain of 8 dBi. The structure of the assumed 64T64R antenna array is depicted in Figure 3, characterised by subarrays of two co-polarised antenna elements. The 16T16R antenna array has full-column subarrays of eight co-polarised antenna elements and hence supports only horizontal beamforming. The assumed maximum transmission power of a cell is 40 dBm (10 Watt).

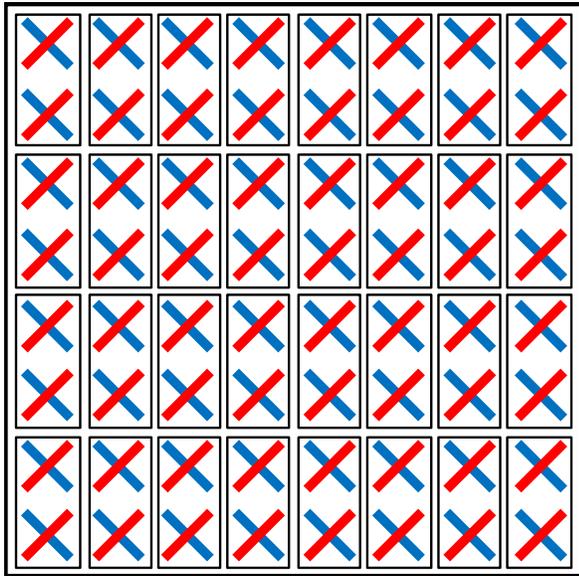


Figure 3: The considered 64T64R antenna array comprises 8 × 8 × 2 cross-polarised antenna elements, grouped in vertical subarrays of 2 × 1 co-polarised antenna elements.

We assume a C-RAN (Cloud/Centralised Radio Access Network) architecture with a functional split at the physical layer, with the remote radio heads (RRHs) deployed at the cell sites and baseband units (BBUs) aggregated in centralised BBU pool. An infinite-capacity fronthaul is assumed between the RRHs and the BBU pool with a 10 ms latency.

Lastly, each cell is assigned a 5 MHz TDD carrier (or slice) in the 3.5 GHz band with the TDD frame configured to support a 4:1 down- versus uplink resource split. We assume use of the 15-kHz numerology, thus providing 25 PRBs (Physical Resource Block) on the assigned carrier.

Regarding *traffic characteristics*, we assume persistent full-buffer data flows maintained for an average of M UEs per cell, whose locations are uniformly sampled within each cell. Denote with $\mathcal{M} = \{1, \dots, 57 \times M\}$ the set of all UEs. It is noted that due to the random nature of shadowing a UE's best serving cell may differ from the nearest one. Each UE is equipped with two cross-polarised antennas, denoting $N_R = 2$.

The assessment is conducted for two distinct *propagation environments*, viz. the urban macro (UMa) and the rural macro (RMa) environments. The environment characterisations in terms of path loss, shadowing and multipath fading are taken from [10] with implementations provided by [11]. Denote with

$$\mathbf{H}_{cm}(t) = [\mathbf{h}_{cm,ij}(t)]_{i=1, \dots, N_R, j=1, \dots, N_T} \in \mathbb{C}^{N_R \times N_T}$$

the wideband channel response matrix at time t between cell $c \in \mathcal{C}$ and UE $m \in \mathcal{M}$, with

$$G_{cm} = \frac{1}{T} \int_{t=0}^T \frac{1}{N_R} \sum_{i=1}^{N_R} \left| \sum_{j=1}^{N_T} \mathbf{h}_{cm,ij}(t) \right|^2 dt$$

the time-average propagation gain (T is the simulation time) and with $\mathbf{h}_{cm,i}(t)$ the i -th row vector of $\mathbf{H}_{cm}(t)$. In notation used below, \mathbf{h}^H refers to taking the complex conjugate transpose of complex vector \mathbf{h} .

Table 1: Simulation model parameters.

PARAMETER	URBAN ENVIRONMENT	RURAL ENVIRONMENT
Network size	57 cells	
Inter-site distance	500 m	5000 m
BS height	25 m	35 m
BS antenna array ($\sim N_T$)	16T16R, 64T64R	
BS transmit power	40 dBm	
Carrier frequency	3.5 GHz	
Carrier bandwidth	5 MHz	
# UEs	57 × M	
UE height	1.5 m	
# UE antennas ($\sim N_R$)	2	
Channel model	3GPP UMa	3GPP RMa

3 Radio resource management

In this section we describe the relevant radio resource management schemes included in the assessment.

The *cell assignment* scheme first selects for UE m the best-serving cell $c^* = \operatorname{argmax}_{c \in \mathcal{C}} G_{cm}$ as the cell with the highest average propagation gain to the UE. When using D-MIMO technology, UE m 's cluster \mathcal{C}_m of serving cells is expanded beyond comprising cell c^* only, by adding any other cell whose average propagation gain to the given UE is at most α dB lower than that of c^* :

$$\mathcal{C}_m = \left\{ c \in \mathcal{C} \mid 10 \log_{10} \frac{G_{c^*m}}{G_{cm}} \leq \alpha \right\}$$

where α is a configurable parameter. Note that in the extreme case of $\alpha = 0$ D-MIMO is effectively turned off and each UE will be served a single cell. At the other extreme, for $\alpha \rightarrow \infty$, for each UE the serving cell cluster will comprise all \mathcal{C} cells, i.e. $\mathcal{C}_m = \mathcal{C}$ for all $m \in \mathcal{M}$.

A *packet scheduler* decides for each TTI (Transmission Time Interval) and each cell which UEs are served. When scheduled in a given TTI, a UE is always concurrently served by all cells in its cluster. To govern the scheduling decisions, for each UE $m \in \mathcal{M}$ a proportional fairness [12] ratio $P_m(t)$ is updated in each TTI, given by

$$P_m(t) = \frac{R_m(t)}{\hat{R}_m(t-1)},$$

where $R_m(t)$ denotes the instantaneous bit rate at which UE m can be served in upcoming TTI t , and $\hat{R}_m(t-1)$ denotes the average throughput experienced by UE m up to and including TTI $t-1$. After each TTI t and for each UE m , $\hat{R}_m(t)$ is updated (exponentially smoothed) as follows:

$$\hat{R}_m(t) = \beta R_m(t) \mathbb{1}\{\text{UE } m \text{ is served}\} + (1 - \beta) \hat{R}_m(t-1),$$

noting the indicator function to capture the scheduling decision on whether or not UE m was actually served in TTI t . At the beginning of the simulations, for each UE m , $\hat{R}_m(0)$ is initialised to 0. A typical setting of $\beta = 0.01$ is assumed, in line with e.g. [13].

At a given TTI t , the centralised packet scheduler ranks the UEs by their proportional fairness ratio $P_m(t)$, where the UE with the highest ratio is ranked first. The scheduler then goes through the list one by one and for each UE decides whether it is scheduled, considering the set of higher-ranked UEs that may have been selected already. More concretely, we employ an adaptation of the heuristic *semi-orthogonal user selection* scheme [14] as outlined below:

STEP 0:

INITIALISE scheduled UE set $\mathcal{M}_{SS,c} = \emptyset$ for each cell $c \in \mathcal{C}$ and global candidate UE set $\mathcal{M}_{CS} = \mathcal{M}$

∀ UE $m \in \mathcal{M}$ determine the best-targeted receive antenna $r_m = \operatorname{argmax}_{i=1, \dots, N_R} \mathbf{h} \mathbf{h}^H$ with $\mathbf{h} = [\mathbf{h}_{cm,i}(t)]_{c \in \mathcal{C}_m}^{\text{row}}$ the row-

wise concatenation of channel response vectors $\mathbf{h}_{cm,i}(t)$ for all cells c in UE m 's serving cell cluster \mathcal{C}_m

STEP 1:

WHILE $\mathcal{M}_{CS} \neq \emptyset$ DO

 Consider $m^* = \operatorname{argmax}_{m \in \mathcal{M}_{CS}} P_m(t)$

 IF $\mathbf{h}_{cm^*,r_{m^*}}(t) \mathbf{h}_{cm^*,r_{m^*}}^H(t) \leq \gamma |\mathbf{h}_{cm^*,r_{m^*}}(t)| |\mathbf{h}_{cm^*,r_{m^*}}(t)|$ for all $c \in \mathcal{C}_{m^*}$, $m \in \mathcal{M}_{SS,c}$ with $m \neq m^*$, THEN SCHEDULE UE m^* and hence $\mathcal{M}_{SS,c} = \mathcal{M}_{SS,c} + \{m^*\}$ for each $c \in \mathcal{C}_{m^*}$

$\mathcal{M}_{CS} = \mathcal{M}_{CS} - \{m^*\}$

END

In STEP 0, the best-targeted receive antenna r_m of UE m is selected as the one towards which the Maximum Ratio Transmission (MRT) precoder (see below) can achieve the strongest signal, utilising the transmit antennas of all cells in the UE's serving cell cluster.

The checks conducted in STEP 1 effectively verify that in each cell in a candidate UE's serving cell cluster the channel between the cell's transmit antennas and the UE's best-targeted receive antenna is sufficiently uncorrelated with the channels of those same transmit antennas towards the best-targeted receive antenna of UEs already scheduled in that cell. Herein, 'sufficiently' is characterised by configuration parameter $\gamma \in [0,1]$. Note that for $\gamma = 0$, the inequality is never satisfied and in each cell at most one UE is scheduled (*SU-MIMO*). At the other extreme, for $\gamma = 1$, all UEs end up being scheduled. In between these extremes, the configuration parameter γ can be tuned to optimise the degree of *MU-MIMO* that is enforced.

Once the scheduler has decided which UEs to serve in each cell, including the options of co-scheduling multiple UEs in a given cell (*MU-MIMO*) as well serving a given UE by multiple cells (*D-MIMO*), the next step is to derive suitable *beamforming* precoders for these transmissions. Considering the *D-MIMO* feature, precoders cannot be determined on a per-cell level yet require a more 'global perspective', which benefits from the assumed *C-RAN* architecture. We first define the complex channel response matrix $\mathbf{H}(t) \in \mathbb{C}^{M_S \times (C \times N_T)}$ as

$$\mathbf{H}(t) = \left[[\mathbf{h}_{cm,r_m}(t) \cdot \mathbb{1}\{c \in \mathcal{C}_m\}]_{c \in \mathcal{C}}^{\text{row}} \right]_{m \in \mathcal{M}_{SS}}^{\text{col}}$$

with $\mathcal{M}_{SS} = |\cup_{c \in \mathcal{C}} \mathcal{M}_{SS,c}|$ and $M_S = |\mathcal{M}_{SS}|$ the aggregate set and total number of scheduled UEs, respectively, and 'col' indicating the column-wise concatenation of the given (themselves concatenated) rows, one for each scheduled UE. Note that each row includes the channel response vector of the associated UE towards all cells in its serving cell cluster and vectors of 0's for all non-serving cells. Given $\mathbf{H}(t)$ we apply zero forcing (ZF) [6] to derive the applied precoders, which are included in precoding matrix $\mathbf{W}(t) \in \mathbb{C}^{(C \times N_T) \times M_S}$, given by

$$\mathbf{W}(t) = \mathbf{H}^H(t) (\mathbf{H}(t) \mathbf{H}^H(t))^{-1} = \begin{bmatrix} \mathbf{w}_{1m}(t) \\ \vdots \\ \mathbf{w}_{Cm}(t) \end{bmatrix}_{m \in \mathcal{M}_{SS}}^{\text{row}}.$$

We then extract for each scheduled UE $m \in \mathcal{M}_{SS}$ (only) the precoding (column) vectors associated with all its serving cells $c \in \mathcal{C}_m$, i.e. $\mathbf{w}_m(t) = [\mathbf{w}_{cm}(t)]_{c \in \mathcal{C}_m}^{\text{col}}$ and subsequently normalise the obtained aggregate precoding vector: $\mathbf{w}_m(t) = \mathbf{w}_m(t) / \|\mathbf{w}_m(t)\|$.

Lastly, on a per-cell level, the available transmit power is evenly split over the layers of all co-scheduled UEs. It is noted that for a UE that is served in D-MIMO fashion and is co-scheduled with different numbers of other UEs in distinct cells of its serving cell cluster, this per-cell power sharing may affect the relative amplitudes among the aggregated set of serving antennas and consequently affect the side lobe suppression [2]. The relative phases and hence the directionality of the beams is however unaffected by this.

In the case a given UE is not co-scheduled with any other UEs in any of its serving cells no zero forcing is needed in deriving its optimal precoder and we simply apply MRT-based precoding, known to maximise the gain [15]. In fact, the above-described procedure is readily verified to indeed yield the MRT precoder.

Once the beamforming precoders have been derived, the *adaptive modulation and coding* scheme estimates the SINR for each UE and maps this to a selected MCS (Modulation and Coding Scheme) based on a 10% BLER (Block Error Rate) target and link-level results available from [16]. To adequately compensate for feedback/measurement delay-induced errors in deriving the optimal transmission parameters, a dynamically tuned UE-specific outer-loop link adaptation scheme is applied.

4 Scenarios & KPIs

The key objective of the presented assessment is to conduct a sensitivity analysis of selected performance metrics w.r.t. different settings of two configurable radio resource management parameters, viz. D-MIMO parameter α and MU-MIMO scheduling parameter γ , and doing so for distinct environment, deployment and load scenarios. Concrete scenario parameterisations are summarised in Table 2.

Table 2: Overview of the simulation scenarios.

PARAMETER	VALUE RANGE
Environment	{urban, rural}
BS antenna array	{16T16R, 64T64R}
# UEs	$57 \times M$, with $M \in \{1, 5, 10\}$
α (~ D-MIMO)	{0, 3, 6, ..., 30} dB
γ (~ MU-MIMO)	{0, 0.1, 0.2, ..., 1}

As mentioned before, during the simulations performance statistics are collected only from the 7×3 ‘evaluation cells’. Two KPIs are of primary interest and will be used to present the results, viz. the average user throughput and the 10th user throughput percentile, which is often regarded a suitable metric

capturing the cell edge performance and generally considered as a driving KPI in mobile network planning and optimisation.

5 Numerical results

The modelling aspects and radio resource management schemes detailed in Section 3 and 4 have been implemented in a system-level simulator. In this section we present the numerical results obtained via the simulator for a variety of scenarios.

5.1 Analysis of D/MU-MIMO on throughput performance

In the first set of scenarios we analyse the impact of the D-MIMO parameter α and the MU-MIMO parameter γ on the throughput performance, considering an urban environment and a deployment of 64T64R antennas at all BSs. We assume a high-load scenario with $M = 10$. Consider first the chart showing the *average user throughput* performance at the top of Figure 4. For the case without D-MIMO, i.e. looking along the axis where $\alpha = 0$ dB, observe that the average user throughput is increasing in γ , revealing the gains from MU-MIMO. In fact, the optimal configuration of γ is equal to 1, i.e. allowing all UEs in a cell to be co-scheduled. The higher γ , the more likely a candidate UE passes the correlation check in STEP 2 of the scheduler and hence the higher the number of co-scheduled UEs. The observed gains are the net effect of a number of contradicting effects; the higher the number of co-scheduled UEs: (i) the lower the per-UE transmit power; (ii) the higher the potential intra-cell (inter-UE) interference; consequently, (iii) the lower per-UE SINR; yet (iv) the higher the fraction of TTIs in which any given UE is scheduled for transmission.

Using Shannon-Hartley’s channel capacity theorem [17] it is readily verified that, unless the intra-cell interference is too high, co-scheduling multiple UEs (with accordingly reduced SINRs) indeed yields a higher aggregate cell throughput and hence a higher spectral efficiency than single-UE scheduling, particularly when the UEs are at relatively high-SINR locations. To see this, consider a simple example with two UEs at identical locations experiencing a relatively high SINR of 15 dB when served in a sequential SU-MIMO fashion. Given a carrier bandwidth of, say, 20 MHz, this yields a per-UE throughput of

$$R_{UE,SU-MIMO} = \frac{1}{2} \cdot 20 \cdot \log_2 \left(1 + 10^{\frac{15}{10}} \right) \approx 50.3 \text{ Mb/s}$$

and hence a cell throughput of 100.6 Mb/s. If the UEs are co-scheduled in MU-MIMO fashion and we assume negligible inter-UE interference, the SINRs are halved to 12 dB due to power sharing, the per-UE throughput becomes

$$R_{UE,MU-MIMO} = 20 \cdot \log_2 \left(1 + 10^{\frac{12}{10}} \right) \approx 81.5 \text{ Mb/s}$$

and the cell throughput is equal to 163.0 Mb/s, an increase of about 62%. If we analyse a similar scenario but assume the UEs are less favourably located with a SU-MIMO SINR of e.g. 0 dB, then the per-UE throughputs are 10.0 Mb/s and 11.7 Mb/s for sequential and co-scheduling, respectively, hence yielding respective cell throughputs of 20 Mb/s and 23.4 Mb/s, i.e. a substantially lower multi-user scheduling gain of 17%.

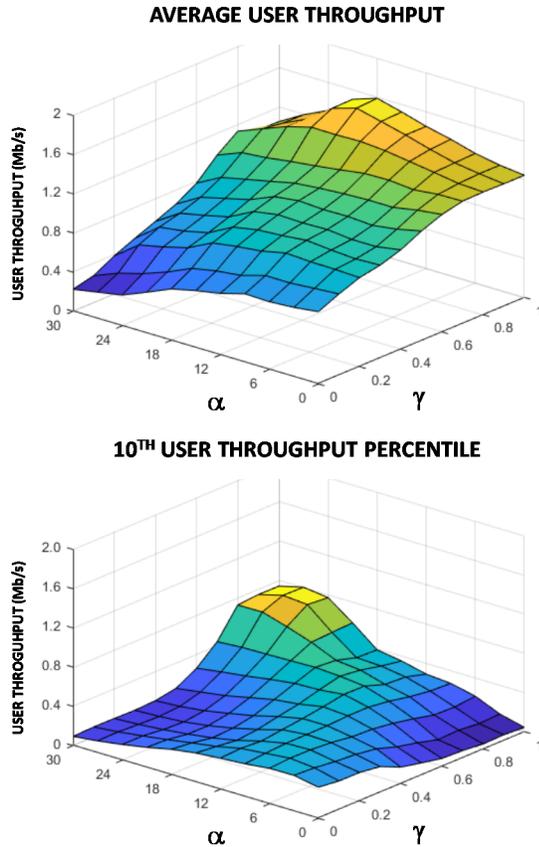


Figure 4: Throughput performance versus D-MIMO parameter α and MU-MIMO parameter γ for the urban environment with a 64T64R antenna deployment and a load of $M = 10$.

In the same chart, consider now the case without MU-MIMO, i.e. looking along the axis where $\gamma = 0$, we observe that the average user throughput decreases in D-MIMO parameter α . In general the cell edge UEs are the most likely candidates to exploit this feature, establish a cluster of serving cells and benefit from the increase in assigned resources in the sense of an enhanced throughput (see below). This more generous resource assignment aiding the cell edge UEs effectively reduces the resource availability for all other users and hence comes at the cost of a consequent reduction in the average user throughput performance, as indeed shown in the chart.

Considering the combined utilisation of both D-MIMO and MU-MIMO, i.e. scenarios with $\alpha, \gamma > 0$, we learn that the use of D-MIMO can indeed also enhance average user throughput, but only in combination with MU-MIMO. As the chart shows, the configuration $(\alpha, \gamma) = (18, 1)$ optimises the average user throughput, yielding an average user throughput of about 1.59 Mb/s. This exceeds the average user throughput that can be attained by using MU-MIMO only (1.25 Mb/s; $\gamma = 1$) or by using neither feature (0.74 Mb/s; $\alpha = \gamma = 0$) by 27% and 115%, respectively. The interesting observation that D-MIMO improves

the average user throughput only in combination with MU-MIMO can be intuitively explained by recalling that the gains from MU-MIMO are most prevalent for high SINRs and it is in fact D-MIMO that enhances the SINRs where they are lowest, i.e. near the cell edge. Worded differently, the use of D-MIMO improves the overall SINR performance by investing more transmit power and beamforming potential in UEs with relatively poor SINRs, while this SINR improvement is subsequently exploited by MU-MIMO to enhance the throughput performance. Moreover, when using MU-MIMO with ZF-based precoding, any significant interference that may exist is inter-cell interference, which is substantially reduced with the use of (also ZF-based) D-MIMO.

Concentrate now on the 10^{th} user throughput percentile (bottom chart in Figure 4), implicitly showing the cell edge performance. The numerical results show that the sole usage of the D-MIMO feature yields a significant improvement of the cell edge throughput. This contrasts with what was seen for the average user throughput but is in line with the very purpose and general potential of D-MIMO. The cell edge throughput is maximised for $\alpha = 3$ dB, yielding a 10^{th} user throughput percentile of 0.36 Mb/s in this high-load scenario, i.e. about 15% higher than that achievable without use of D/MU-MIMO (0.32 Mb/s; $\alpha = \gamma = 0$). Also in contrast with what was observed for the average user throughput, the sole usage of MU-MIMO actually degrades cell edge performance. Although this could be interpreted to suggest that MU-MIMO is a useless feature when optimizing for cell edge performance, the most interesting insights are however obtained when analysing the combined use of D/MU-MIMO. As the chart shows, for each choice of the MU-MIMO parameter γ , raising the D-MIMO parameter α first enhances the cell edge performance up to a certain maximum, beyond which the performance degrades again due to over-investing of resources in D-MIMO operations. Furthermore, the higher the choice of γ , the higher the optimal α , i.e. the more UEs are co-scheduled in MU-MIMO fashion, the higher the degree of D-MIMO that should be applied to enhance the ‘worst SINRs’ and maximise the joint D/MU-MIMO performance gains. Clearly, a well-tuned combination of both features yields the highest performance. More specifically, configuration $(\alpha, \gamma) = (27, 1)$ optimises the cell edge performance, yielding a 10^{th} user throughput percentile that can be attained by using D-MIMO only (0.36 Mb/s; $\alpha = 3$ dB) or by using neither feature (0.32 Mb/s; $\alpha = \gamma = 0$) by 121% and 153%, respectively. Observe that the best joint configuration of both features is different when optimizing for average of cell edge performance.

5.2 Performance impact of the traffic load, propagation environment and the deployed antenna array

The scenarios analysed above considered a fixed high-load scenario with $M = 10$ and only addressed the deployment of 64T64R antennas in an urban environment. In this subsection we present a sensitivity analysis of the optimised cell edge performance w.r.t. the traffic load, the propagation environment and the deployed antenna arrays. The corresponding results are

shown in Figure 5 for the urban (top chart) and rural environments (bottom chart). Specifically, each chart shows both the average throughput and the 10th throughput percentile for a low-load ($M = 1$), a medium-load ($M = 5$) and a high-load ($M = 10$) scenario and for antenna array options 16T16R and 64T64R. For each scenario the throughput metrics are shown for the best possible configuration of (α, γ) , where the optimisation is done w.r.t. the 10th user throughput percentile.

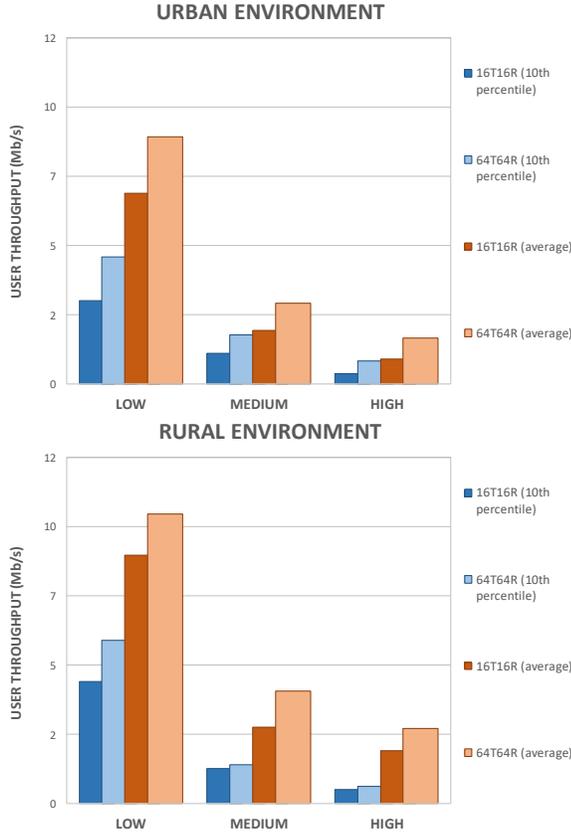


Figure 5: Throughput performance for the urban and rural environments with 16T16R or 64T64R antenna deployments and a low, medium or high load, assuming optimised D/MU-MIMO configurations.

Aside from the more evident observation that all throughput curves decrease in the traffic load due to increased competition for resources, a few more insightful observations can be made from the obtained results. Firstly, we observe that the cell edge throughput gain from having 64T64R over 16T16R antennas is largest in the urban environment. This is due to the fact that the network deployment in the urban environment is characterised by a smaller inter-site distance, hence UE-BS distances are smaller, implying a higher potential for vertical beamforming, which is supported only by the 64T64R antenna. For the rural environment the gain from having 64T64R antennas is relatively small and unlikely to be worth the additional investment cost.

Although not explicitly shown in the figure, we further note that in all scenarios the optimal configuration of the MU-MIMO

parameter γ is equal to 1, hence in each cell all available UEs are co-scheduled, while the optimal setting of the D-MIMO parameter α is higher for the case with 64T64R than for the case with 16T16R antennas. A possible reason for this observation is that 64T64R antennas have better beamforming capabilities and hence also inherently yield better MU-MIMO performance, implying that larger serving cell clusters may be needed to utilise the full D-MIMO potential. Lastly, we learn from the underlying numerical results that the performance gains from D-MIMO are higher in case of 16T16R deployment than for a 64T64R deployment, which is likely due to the fact that in the latter case the number of antennas used for beamforming is already relatively large, while in the former case the number of beamforming antennas can be more substantially increased beyond that of a single-cell case by involving antennas from multiple cells in D-MIMO fashion.

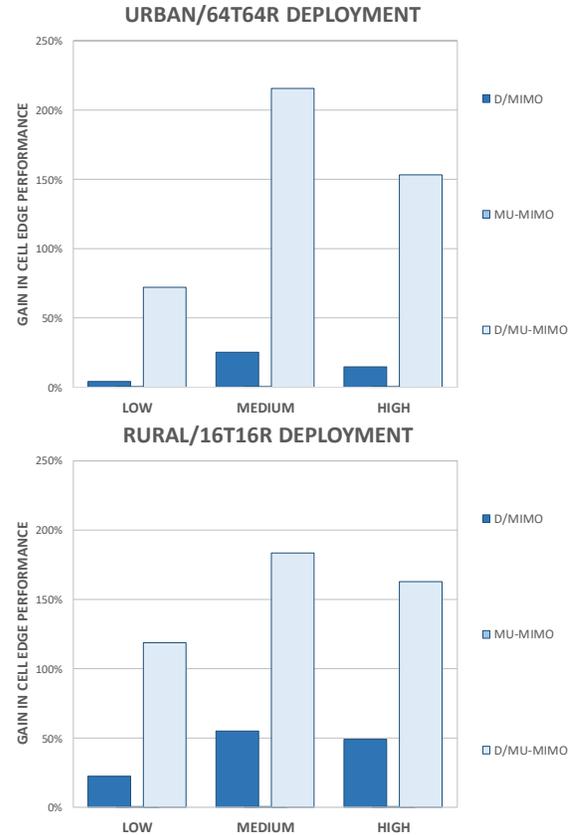


Figure 6: Gain in the 10th user throughput percentile from optimised D-MIMO, MU-MIMO or D/MU-MIMO configurations for the urban (rural) environment with a 64T64R (16T16R) antenna deployment and a low, medium or high load.

5.3 Attainable performance gains from D-MIMO, MU-MIMO and combined D/MU-MIMO

The analyses presented in Section 5.1 already revealed for a specific scenario that it is the combined use of D-MIMO and MU-MIMO with jointly optimised configurations that yields the

highest throughput performance. In this section, we further analyse the attainable percentual gains in the cell edge performance from either feature in isolation as well as for their combined use, with reference to the case where neither feature is used. The results are shown in the bar charts of Figure 6, for low ($M = 1$), medium ($M = 5$) and high load scenario ($M = 10$). The top (bottom) chart considers the urban (rural) environment with 64T64R (16T16R) antenna array deployments in each cell, considering the practically most likely deployment choices in light of the performance-vs-cost trade-off.

The charts shown in the figure reveal that for both considered environments/deployments there are no gains in the cell edge performance from using only MU-MIMO, as observed before. The gains from exclusive use of D-MIMO are substantial, while the gains from joint use of the D-MIMO and MU-MIMO features are even more significant. In both scenarios, the percentual gains appear to be highest for medium traffic loads.

6 Concluding remarks

In this paper we presented a simulation-based assessment of the performance potential of D-MIMO, MU-MIMO and, in particular, the joint usage of both performance-enhancing features. In the process we have extended previously published scheduling and beamforming principles to enable support of the combined D/MU-MIMO operation. Key insights revealed by the assessment study include the observation that, when optimizing average user throughput performance, D-MIMO, while fruitless when used in isolation, is very effective when intelligently combined with MU-MIMO. Alternatively, when optimizing the cell edge throughput, MU-MIMO, while also ineffective when used in isolation, is in fact very valuable when accompanied by a suitable configuration of D-MIMO. As an illustrative example, when using optimally configured joint use of D/MU-MIMO in a highly loaded urban deployment scenario, a 121% cell edge performance gain can be attained over a scenario using only D-MIMO, and even a demonstrated 153% gain over a scenario where only the MU-MIMO feature is available.

In our continued research we intend to compare the demonstrated performance enhancements with those achieved in

different scenarios, including scenarios with non-persistent traffic flows, Grid of Beams-based beamforming and investigating the impact of load and latency in the fronthaul of the C-RAN architecture.

REFERENCES

- [1] H. Holma, A. Toskala and T. Nakamura, '5G technology - 3GPP New Radio', John Wiley & Sons, Hoboken, USA, 2020.
- [2] M. Kottkamp, A. Pandey, D. Raddino, A. Roessler and R. Stuhlfauth, '5G New Radio - Fundamentals, procedures, testing aspects', Rohde & Schwarz, Munich, Germany, 2019.
- [3] S. Venkatesan, A. Lozano, R. Valenzuela, 'Network MIMO: overcoming intercell interference in indoor wireless systems', *Proceedings of Asilomar Conference on Signals, Systems and Computers '07*, Pacific Grove, USA, 2007.
- [4] E. Dahlman, S. Parkvall and J. Sköld, '4G: LTE/LTE-Advanced for Mobile Broadband', 2nd edition, Academic Press, Oxford, United Kingdom, 2014.
- [5] G. Interdonato, E. Björnson, H. Quoc Ngo, P. Frenger and E.G. Larsson, 'Ubiquitous cell-free Massive MIMO communications', *EURASIP Journal on Wireless Communications and Networking*, 2019.
- [6] A. Davydov, G. Morozov, I. Bolotin and A. Papathanassiou, 'Evaluation of joint transmission CoMP in C-RAN based LTE-A HetNets with large coordination areas', *Proceedings of Globecom '13 workshops*, Atlanta, USA, 2013.
- [7] V. Jungnickel, K. Manolakis, W. Zirwas, B. Panzner, V. Braun, M. Lossow, M. Sternad, R. Apelfrojd and T. Svensson, 'The role of small cells, coordinated multipoint and massive MIMO in 5G', *IEEE Communications Magazine*, vol. 52, no. 5, 2014.
- [8] J. Jin, Q. Wang, C. Lin, H. Yang and Y. Wang, 'Coordinated multi-point transmission with limited feedback', *Proceedings of Globecom '10*, Miami, USA, 2010.
- [9] G. Wang and M. Lei, 'Enabling downlink coordinated multi-point transmission in TDD heterogeneous network', *Proceedings of VTC '13 (Spring)*, Dresden, Germany, 2013.
- [10] 3GPP, 'Study on channel model for frequencies from 0.5 to 100 GHz', TR38.901, v16.1.0, 2019.
- [11] Fraunhofer, 'Quadriga - The next generation radio channel model', www.quadriga-channel-model.de, 2020.
- [12] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushyana and A. Viterbi, 'CDMA/HDR: a bandwidth efficient high speed wireless data service for nomadic users', *IEEE Communications Magazine*, vol. 38, no. 7, 2000.
- [13] A. Jalali, R. Padovani and R. Pankaj, 'Data throughput of CDMA-HDR a high efficiency-high data rate personal communication wireless system', *Proceedings of VTC '00 (Spring)*, Tokyo, Japan, 2000.
- [14] T. Yoo and A. Goldsmith, 'On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming', *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 3, pp. 528-541, 2006.
- [15] E. Björnson, M. Bengtsson and B. Ottersten, 'Optimal multiuser transmit beamforming: a difficult problem with a simple solution structure', *IEEE Signal Processing Magazine*, vol. 31, no. 4, pp. 142-148, 2014.
- [16] S. Pratschner, B. Tahir, L. Marijanovic, M. Mussbah, K. Kirev, R. Nissel, S. Schwarz and Markus Rupp, 'Versatile mobile communications simulation: the Vienna 5G link level simulator', *EURASIP Journal on Wireless Communications and Networking*, 2018.
- [17] C. E. Shannon, 'Communication in the presence of noise', *Proceedings of the Institute of Radio Engineers*, vol. 37, no. 1, 1949.