Field Trials using Coordinated Multi-Point Transmission in the Downlink

V. Jungnickel[+], A. Forck[+], S. Jaeckel[+], F. Bauermeister[+], S. Schiffermueller[+], S. Schubert[+]
S. Wahls[+], L. Thiele[+], T. Haustein[+], W. Kreher[*], J. Mueller[*], H. Droste[*], G. Kadel[*]
[+] Fraunhofer Heinrich-Hertz-Institute, Einsteinufer 37, 10587 Berlin, Germany
[*] Deutsche Telekom AG, Deutsche-Telekom-Allee 7, 64295 Darmstadt, Germany

***Abstract*—We report on field trials using CoMP transmission in the downlink of a mobile radio network. Two new features enable over-the-air CoMP transmission from physically separated base stations and terminals. These are distributed synchronization and a fast virtual local area network. Using VLAN tags, terminals feed back the multi-cell channel state information to their serving bases where it is multiplexed with shared data. Both are multicast to other cooperative base stations over the backhaul. In our trials, two terminals are served in two overlapping cells and placed in specific indoor, outdoor-to-indoor and outdoor scenarios. We have realized both intra-site as well as inter-site CoMP. While outage is indeed a big problem at the cell edge with full frequency reuse, with CoMP it is not observed anymore. Average throughput gains by factors 4 to 22 are observed when using CoMP compared to interference-limited transmission while between 27 and 78% of the isolated cell throughput is measured in both cells simultaneously.**

I. Introduction

Inter-cell interference is the major bottleneck for the performance of mobile networks nowadays. If terminals are served on the same radio resource in adjacent cells they experience severe co-channel interference. At the cell edge, normally the signal of the serving base station (BS) is received with similar power compared to the sum of the signals received from other BSs. One way to cope with this critical situation is coordinated scheduling (CoSCH), where the terminals provide feedback about their possible signal to interference and noise ratio (SINR) for a set of designated scheduling decisions in adjacent cells. Now, the serving BS negotiates with adjacent BSs in which cell which resource is preferably assigned. In a wider sense, CoSCH is a closed-loop interference-aware implementation of the classical frequency reuse [1]. As a consequence, resources remain unused in some cells.

Recent research has shown that it is more efficient to share resources even at the cell edge provided that there is a high-speed backbone between the BSs over which information can be exchanged with negligible latency [2–10]. Joint transmission can be considered as a distributed multiple-input multiple-output (MIMO) system where multiple BSs perform joint beam-forming for multiple terminals served in multiple cells. This scheme is denoted as network MIMO [9] and as coordinated multi-point (CoMP) in recent standardization documents [11].

Ideas for a distributed implementation are reported in [12–17], see Fig. 1. BSs are synchronized using the global positioning system (GPS). Terminals estimate the channel state

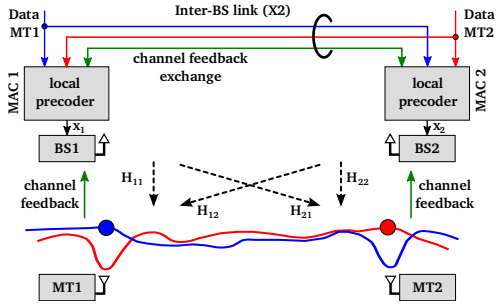

Figure 1. Distributed coordinated multi-point transmission.

information (CSI) using cell-specific reference signals and deliver CSI feedback to their serving BS over the reverse link. Cooperating BSs exchange CSI as well as shared data over a low-latency signaling network denoted as X2 [18]. Weights for the joint beam-forming are computed redundantly at each BS. The locally relevant set of weights is applied to the data signals to calculate the transmitted waveforms. Over the air, the desired signals sum up constructively while the mutual interference between the cells is cancelled.

We have implemented the distributed CoMP concept in real time on top of an existing LTE trial system described previously [19–21]. Implementation of enabling features such as clock synchronization, synchronous data exchange, cell specific pilots, CSI feedback, cooperative precoding, precoded pilots and initial laboratory trials are reported in [22]. Here we describe further system extensions enabling over-the-air experiments with physically separated BSs and distributed mobile terminals (MTs) and report initial field trial results.

Our paper is organized as follows. We describe further enabling features in Section II. In Section III, transmitter and receiver configurations in the field are depicted. Measured quantities are defined in section IV. Measurement results are reported in Section V.

II. Further Enabling Features

Distributed implementation is the key for introducing CoMP in next generation mobile networks. In this chapter, we explain how we realize phase-coherent transmission of the base stations at physically separated sites without synchronizing them by wire. Secondly, we explain how we organize CSI feedback and communication between the base stations using low-cost Ethernet network equipment.

978-1-4244-9116-2/10/$26.00 ©2010 IEEE

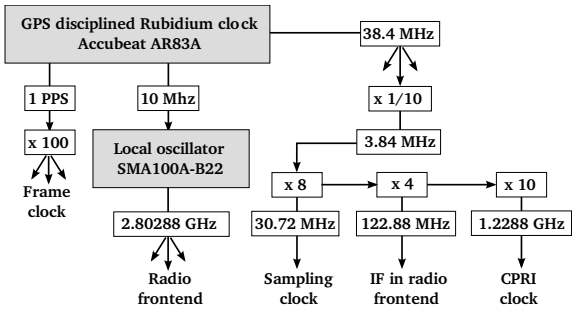

Figure 2. Synchronization and clock distribution in our trial system.

A. Distributed synchronization

For stable interference suppression we need phase-coherent transmitters. As a reference clock at each BS site, we have used a GPS disciplined clock, see Fig. 2. Our customized AR83A has the following output signals locked to GPS: 1 pulse per second (PPS) clock (3x), 10 MHz reference (1x), 38.4 MHz reference (3x) and a serial interface providing location, time and other information in the common National Marine Electronics Association (NMEA) format (3x). The 10 MHz reference feeds a Rhode&Schwarz SMA100A local oscillator (LO) operated at 2.80288 GHz. Clock signals are created once at each site and shared between the sectors. The 38.4 MHz clock fulfills many purposes in the LTE signal processing unit (LSU). Dividing it by 10 results in the 3.84 Hz base clock of the Universal Mobile Telecommunications System (UMTS). Multiplying this base clock by 8, the 30.72 MHz sampling clock for the entire real-time signal processing chain is derived. Multiplication of 30.72 MHz by 4 results in the 122.88 MHz clock which is further multiplied by 10 in a parallel-to-serial converter to obtain the 1.2288 GHz clock for the common public radio interface (CPRI) between LSU and each radio frequency (RF) front-end. The CPRI line receiver recovers the 122.88 MHz clock used also as an intermediate frequency (IF). In the RF front-ends, the LO is mixed with the IF and the difference is our 2.68 GHz carrier frequency used in the downlink.

Note that phase noise becomes crucial if the delay between multi-cell channel estimation and the application of the corresponding precoder weights is 20 ms as in our implementation. Even if we assume a static channel, uncorrelated phase noise at the transmitters effects the interference suppression. Assuming that the precoding matrix $\mathbf{P}$ is the right-handed pseudo-inverse of the channel matrix $\mathbf{H}$, individual phase variations at the distributed transmitters can be described by a phase matrix $\mathbf{\Phi}$ placed between the precoder and the channel. If we consider the 2x2 CoMP setup in Fig. 1 and set the random phase noise $\varphi_1(\tau) \neq \varphi_2(\tau)$, the residual interference is unpredictable

$$\begin{pmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{pmatrix} \cdot \begin{pmatrix} e^{j\varphi_1(\tau)} & 0 \\ 0 & e^{j\varphi_2(\tau)} \end{pmatrix} \cdot \begin{pmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{pmatrix}$$

while for $\varphi_1(\tau) = \varphi_2(\tau)$, the common phase variation can be tracked at the receiver. Therefore we have stabilized the LO phases during the precoder delay using a common low-phase-noise LO for all sectors locked to the GPS at each site.

A puzzling phase noise contribution has been identified in the digital linearization unit for the high-power amplifier in our RF front-ends. Most likely it is due to imperfect calibration of the non-linear gain characteristics of the amplifier. For regular LTE transmission, this unusual contribution is corrected by the phase tracking at the MT while being destructive for CoMP. Since it is not effective at 10 W per-antenna output power, linearization has been switched off.

Note that MTs irregularly correct their carrier frequency and timing offsets and apply individual automatic gain control (AGC) settings at each antenna depending on the received signal strengths. For the computation of the zero-forcing precoder matrix $\mathbf{P}$ we have ignored these terminal settings. For intuition, the zeros in space are steered at the locations of the terminal antennas, irrespective of these settings. Advanced algorithms, such as channel prediction, may need to know these settings, hence they should be reported to the base station in addition to the CSI.

B. CSI and data exchange network

Our signaling network relies on widely available low-cost equipment based on the IEEE 802.3 Ethernet standard. As a general approach, we take profit of an existing standard extension, namely virtual local area networks (VLANs) described in IEEE 802.1q. VLAN enabled switches are used to multiplex and demultiplex the different traffic types for the precoding such as CSI, shared data between the cooperative BSs and internet protocol (IP) based application data of the MTs.

Note that the network architecture depicted in Fig. 3 turned out as an essential enabler for CoMP trials in the field where we use distributed MT and BS locations. In our previous trials [22], the CSI feedback from both MTs has been decoded at each BS. Note that channel gains and delays between the distributed BSs and MTs in the field can vary substantially.

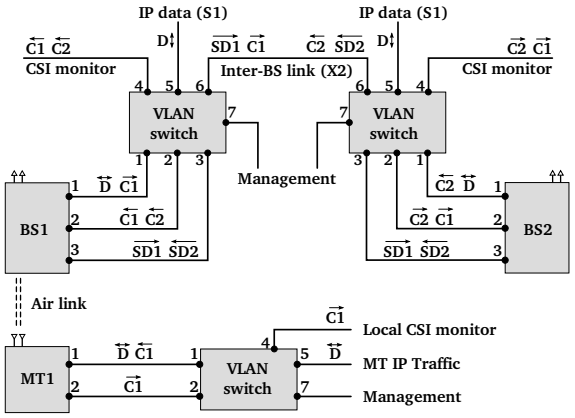

Figure 3. Setup of the virtual local area network (VLAN) for the CSI feedback and the data exchange between the base stations.

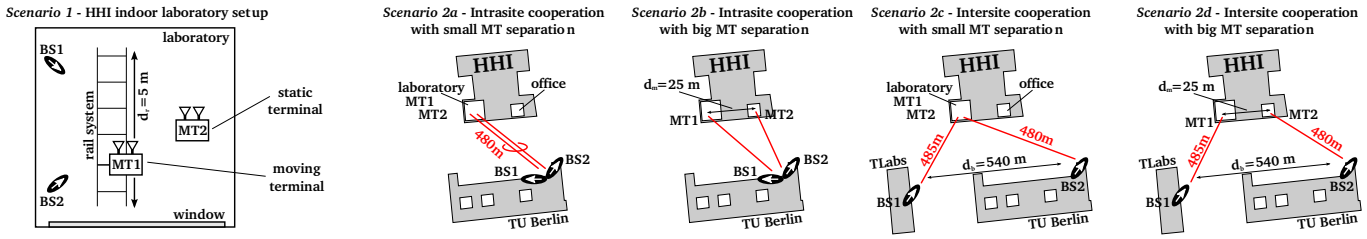

Figure 4. Measurement scenarios 1 and 2. Left: Indoor scenario. Right: Outdoor-to-indoor scenario.

However, since power control and timing advance are steered from the serving base only, the feedback is frequently not decoded correctly at the other BS. Now only the feedback of the served MT is decoded (in the optimized up-link) and the BSs exchange the CSI as shown in principle and in detail in Figs. 1 and 3, respectively. In this way, the CSI feedback becomes quite robust also for distributed setups in the field.

Once in each 10 ms radio frame, a particular digital signal processor (DSP) at the MT takes over the raw estimates of the multi-cell channel matrix from the physical layer implemented in a field programmable gate array (FPGA). The DSP packs the CSI into a standard Ethernet packet using a multicast address. Packets denoted as C1 in Fig. 3 are further processed by the FPGA where they are extracted out of the physical layer processing chain and sent to port 2 of the MT over 100 Mbit/s Ethernet. Next the switch adds a VLAN tag to the C1 packets identifying them as CSI. By using another VLAN tag, IP application data (denoted as D) are identified at input 5. In the VLAN switch, the CSI is multiplexed with data on output 1. Port 1 at the MT can be regarded as a transparent Ethernet tunnel to the serving BS over the wireless link.

At the BS, the multiplexed up-link traffic is fed again into a VLAN switch. Based on the tags, the switch splits CSI and data packets and forwards them into different ports. The CSI stream C1 is copied into several output ports: The input of the precoder at BS1 (2), a mirror of port 2 used for monitoring and measurements (4) and the inter-BS link (6) denoted as X2 interface in the LTE/SAE standardization. In addition to the C1 stream, a shared data stream SD1 containing the already encoded and scheduled downlink data for MT1 is also forwarded to port 6 of the switch. These data are taken out of the signal processing pipeline of the BS immediately after the decentralized MAC layer processing, see [22], Fig. 5. The same Ethernet stream processing is done simultaneously at BS2. As a result, the data streams C2 and SD2 are arriving on port 6 of the switch and forwarded to port 2 and 3 of BS1, respectively. Each BS is now aware of the data of the other BS as well as of the CSI from the MTs in both cells. A particular DSP computes the precoding matrices independently at each BS, transfers the results to an FPGA where the locally relevant weights are applied to the jointly transmitted data. With the resulting signals, both MTs are served on the same resources in both cells while mutual interference is cancelled.

The use of VLAN tags has a negligible impact on the overall feedback delay. The delay between channel estimation and application of the precoder weights is primarily due to other contributions. To illustrate this, we report cumulative delays with respect to the beginning of a radio frame. Coarse multi-cell channel estimation takes 0.3 ms in the downlink receiver. Transferring the results to the DSP is finished after 1.7 ms. Forming CSI packets and delivering them into the uplink is completed after 5 ms. Transferring our uncompressed CSI at a data rate of 4.6 Mbit/s over the uplink is finished after 12 ms. Channel interpolation, computation of the precoding matrices for 1.200 subcarriers and transfer to the FPGA are completed after 19 ms. Note that there is some potential in the feedback chain for reducing the delay: Faster interfaces between signal processing and Ethernet, efficient feedback compression and faster signal processing at MT and BS.

III. Setup and Scenarios

Scenarios comprise specific indoor, outdoor-to-indoor, and outdoor configurations see Figs. 4 and 5. Indoors (scenario 1), both BSs are located in the same lab. In the outdoor-to-indoor and outdoor scenarios, we transmit from two BS sites and select two sectors either at one site or at two sites in order to realize intra- or inter-site cooperation (scenarios 2a-d, 3a-b), respectively. Sites are located on the Deutsche Telekom Laboratories (TLabs) building at Ernst-Reuter-Platz (84 m antenna height) and on the Technical University of Berlin (TUB) main building, Straße des 17. Juni (43 m, see Fig. 5). The estimated height of buildings in the area is in between 25 and 35 m. For more insights, refer to [21], Fig. 6 bottom right. Sites are interconnected by an optical fiber deployed in the campus with a length of 4.5 km. Signaling is based on 1 Gbit/s Ethernet where data and CSI exchange consume 300 and 4.6 Mbit/s, respectively.

For the indoor and outdoor-to-indoor scenarios, both MTs are located on the 11^{th} floor at the Heinrich Hertz Institute (HHI). We have placed both MTs at the south front of the building with the window facing both base stations either in the same lab (scenarios 2a, c) or in two different labs which are 25 m separated (scenarios 2b, d). MT2 is at a fixed location. In order to capture the local fading statistics, MT1 moves at low speed of about 3 cm/s on 5 m long rails in the lab (see Fig. 4, left). In the field scenario, MT2 is at a fixed position at the ground in a van in front of the south side of the HHI building or in the 11^{th} floor lab inside the HHI building. MT1 is moved in a second van to fixed locations in the field indicated in Fig. 5. The assignment of BS1/2 and MT1/2 to the actual locations is available in Figs. 4 and 5. In our implementation, MT1 is always assigned to BS1 and MT2 to BS2, i.e. handover is not performed.

IV. Measured quantities

In the following we explain what quantities we have measured. Joint precoding enhances the SINR of the data streams dedicated for the MTs. In the experiments we use a spatially multiplexed transmission of four data streams from two base stations, i.e. two streams from each. On top of the physical layer, frequency-selective link adaptation and MIMO mode switching is implemented. MTs estimate the effective channel $\mathbf{HP}$ using the demodulation reference signals (DRS). We compute the frequency-selective SINR for 16 groups where each group covers 75 subcarriers. SINRs are computed for a comb of three subcarriers in each group for several transmission modes: single stream transmission on either stream 1 or 2 and transmission of both streams in case of spatial multiplexing. The SINR values for all three options are quantized in four steps (off, QPSK, 16-QAM, 64-QAM) and fed back as a compound frequency-selective channel quality identifier (CQI) vector to the serving BS. Among the three subcarriers per group, the lowest SINR value has been selected. For the SINR-to-rate mapping function we have used a target of 10^{-1} for the uncoded bit error rate.

The BS compares the achievable rates for all spatial transmission modes in the cell. The rate is maximized for each subcarrier group selectively by choosing the best transmission mode and assigning the modulation format recommended by the terminal for each stream and each group selectively. In the sum over the whole bandwidth, a variable data rate is hence assigned.

All data bits transmitted in a transmission time interval (TTI) are considered as a transport block in LTE. The block is encoded using a fixed code rate of 1/2 in our setup and then passed through a flexible interleaver supporting a variable codeword length. Actually measured error rates at the decoder input are in the order of 10^{-2}, due to the conservative subcarrier selection described above. This yields almost error-free transmission after decoding, i.e. only few retransmissions are needed in the hybrid automatic repeat request (HARQ) module. The gross data rate is measured at the physical layer as the sum of the successfully transmitted data and parity bits per second excluding overhead for synchronization, multi-cell channel estimation and control signaling. The peak data rate of our system is 141 Mbit/s in one cell. Throughputs, multi-cell channel coefficients and bit error rates are simultaneously recorded during all measurements.

V. Measurement Results

Indoor and outdoor-to-indoor results are plotted in Fig. 6. Left in each row of figures we have plotted the SINR per data stream in the interference-limited case, i.e. without using CoMP, and after applying optimum combining (OC) based on the DRS at the MT receiver. On the right side, the throughput at each MT is plotted for three transmission modes. As an upper bound, we consider an isolated cell where the interference from the other cell is switched off (blue). As a lower bound, we realize interference-limited transmission using an identity matrix as the precoder in combination with

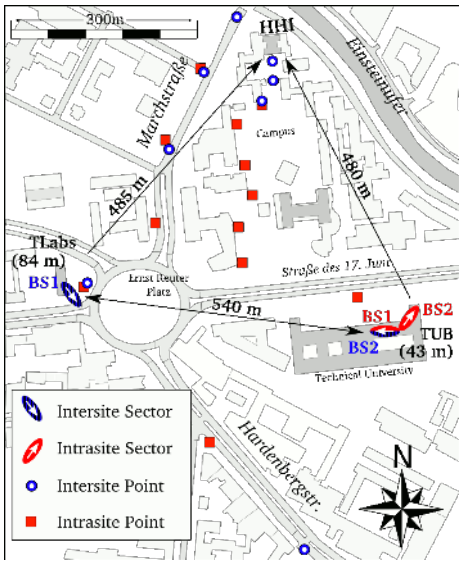

Figure 5. Location of base stations and MT1 in the field trials.

OC at the receiver. Finally, we have used CoMP with a fixed number of four streams (black).

In the indoor scenario 1, both BSs are received with the same average power. This is obvious also from the SINR being around 0 dB in both cells simultaneously. Due to multiple reflections in the room, however, both the signal and the interference experience fading. Independent fading of both components creates a crucial throughput situation for a MT: When moving the MT by a few cm only, we can realize situations where either the signal channel is flat while the interference is in a fade, and correspondingly the serving BS assigns data transmission only in a fraction of the whole frequency band, as well as the reverse situation where the interference channel is flat while the signal is in a fade so that no more data are usually transmitted. As a result, the terminal suffers outage in 50% of the cases, and in general the user experience is poor when moving through the lab.

If CoMP is enabled in such a bad interference scenario, we observe dramatic improvements of the throughput. Despite the critical interference situation and although the data rate still varies, CoMP removes outage at the cell edge completely. Using CoMP in scenario 1, both terminals have on average a 18.7 times higher throughput and realize 78% of the rate in the isolated cell (see Tab. I).

Next consider the intra-site scenarios 2a and 2b. It is typical in the distributed multi-cell network that the path losses are not equal for different pairs of BSs and MTs. Nonetheless, the principal observations remain similar. In the interference-limited case, again there is significant outage and irregularly a few percent of the peak rate can be realized. Using CoMP, in contrast, both MTs can realize 34 and 27% of the peak rate on average in scenarios 2a and 2b, respectively.

The performance in the inter-site scenarios is superior compared to the intra-site scenarios, despite distributed synchronization of both base stations. Determining factors for the CoMP gains are the different SINR situations of the terminals and also the physical separation between the BS as

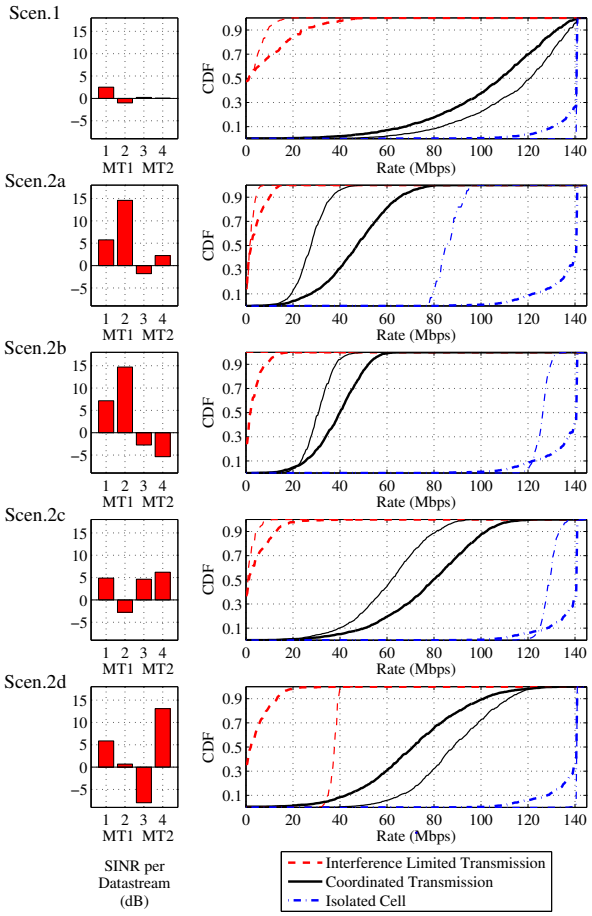

Figure 6. Measurement results for scenarios 1 and 2. Left: SINR per data stream. Right: Measured throughput for interference-limited transmission (red), coordinated multi-point (black) and isolated cells (blue). Thick lines MT1, thin lines MT2.

illustrated now. In previous channel measurement results [23], we observed that the transmitter correlation depends on the antenna spacing between the BSs. In Fig. 7, we have verified this using the recorded multi-cell channels. If we use different polarizations at the BSs and compare the channels to the same MT antenna, little correlation is found. The correlation is higher but still rather low, if we consider the same polarization and both signals come from different sites. But if these co-polarized signals come from the same site, the correlation is significant. Hence, the higher data rates observed with CoMP in the inter-site scenarios may also be attributed to the lower transmit antenna correlation.

Let us finally consider the outdoor results. We have selected only locations where both MTs are served by their strongest BS to consider the handover which will be performed in a real network. The geometry factor (see Fig. 8, top) shows the ratio of the mean signal power to the mean interference power for both, the intra- and inter-site scenarios. The fixed MT2 position at the ground of the HHI building has a GF of -1.2 dB for the intra-site and -4.5 dB for inter-site scenario. Likewise, the other MT2 position at the 11^{th} floor achieved -1.7 dB and +7 dB. We have put all intra- and inter-site locations of MT1 into separate statistics. Obviously, the geometry of the intra-site scenario is too optimistic, with higher SINR compared to typical cellular deployments, while the inter-site scenario contains a fairly realistic interference distribution similar to multicell simulations [24].

Table I
MEAN RATES FOR ALL SCENARIOS [MBIT/S]

Scen.	Int. limted	CoMP	Isolated cell
1	5.8	109.3	139.6
2a	3.1	37.5	111.1
2b	1.7	35.3	131.2
2c	3.1	69.5	133.3
2d	21.1	80.1	139.1
3a	11.0	48.1	104.2
3b	2.1	39.1	98.9

According to the geometry, in the intra-site trials there are several locations where a relatively high data rate is assigned already in the interference-limited case because the scheduler is aware of the SINR which includes both, the channels to the own BS and to the interfering BS (see Fig. 8, bottom). In the inter-site trials, the geometry is more realistic and much lower data rates are assigned. In both scenarios, there is a certain probability of outage also in the field. If we use CoMP, these obvious differences in the geometries do not reflect on the data rates. The observed rates have a similar distribution, possibly due to impairments such as residual time variance of the channel. On average, 46 and 40% of the isolated cell capacity are measured in both cells simultaneously, owing to the active mutual interference suppression using CoMP (see Tab. I).

CONCLUSIONS AND DISCUSSION

In this paper, we have reported for the first time on field trials in a real multicellular deployment using coordinated multi-point (CoMP) transmission in the downlink. Coherent interference nulling has been demonstrated over the air while a number of essential network operator requirements have been met. We have used distributed synchronization and linked base stations using standard Ethernet. The high latency requirements for the information exchange are met using commercially available network equipment based on the IEEE 802.1q virtual local area network standard. In our inter-site scenarios, we demonstrated for the first time downlink CoMP over 20 MHz bandwidth at 500 m inter-site distance. This is the proof of concept that downlink CoMP can be integrated into the distributed LTE system architecture.

We have performed indoor, outdoor-to-indoor and outdoor measurements in the field and demonstrated that the advan-

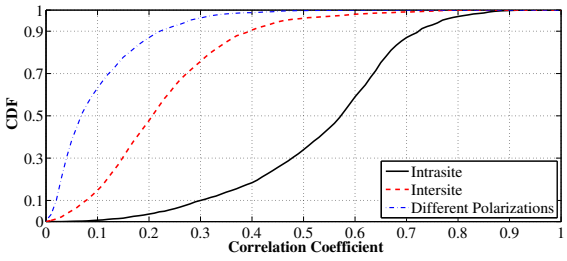

Figure 7. Fading correlation for several transmitter configurations.

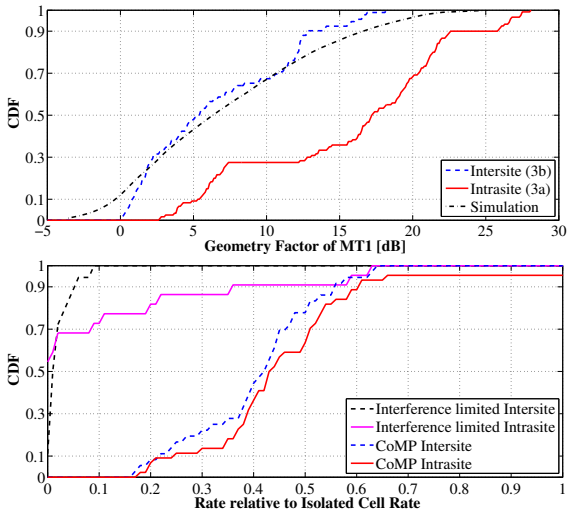

Figure 8. Top: Geometry for MT2 in the field trials. Simulations have been done according to [24] with 10° tilted antennas. Bottom: Throughput compared to isolated cell.

tages of CoMP are indeed remarkable. The gains for the average data rate vary between factors 4 and 22 depending on the scenario. We have turned the unpredictable on-off characteristics of the interference channel next to the cell-edge into a stable continuous data link with some residual rate variation qualifying for advanced real-time multimedia applications such as mobile video conferencing. We have realized on average between 27% and 78% of the isolated cell throughput in the investigated scenarios despite reusing the full frequency band in both cells. It is a particular observation that the resulting throughput variation seems to depend not as much on the original user geometry in the field. The large measured gains illustrate that downlink CoMP can be implemented with high precision and that the implementation challenges can be overcome. But of course, such high gains are hardly realistic for large-scale mobile networks if two cooperative base stations are surrounded by non-cooperating cells. External interference has not been present in our trials and it would lower the gains. Moreover, the performance is reduced if the terminal is mobile due to the precoding delay. The zeros for the interference are then realized at positions where the terminal has been situated earlier. Channel prediction would help a lot, yet it is complicated by timing and frequency corrections irregularly performed at the terminals. In order to make the promising CoMP technology mature for next generation mobile networks, we need further insights into efficient user grouping and dynamic clustering of cooperation areas and into the impairments due to the channel time variance.

Acknowledgements

The German Ministry of Education and Research (BMBF) is acknowledged for support in the project EASY-C under contract number 01BU0631. The authors wish to thank their project partners for the fruitful cooperation in this collaborative research project.

References

[1] P. Frank, A. Müller, H. Droste, and J. Speidel, "Cooperative interference-aware joint scheduling for the 3GPP LTE uplink," *Proc. IEEE PIMRC' 10*, 2010.

[2] P. Baier, M. Meurer, T. Weber, and H. Troger, "Joint transmission (JT), an alternative rationale for the downlink of time division CDMA using multi-element transmit antennas," *Proc. IEEE ISSSTA '00*, 2000.

[3] S. Shamai and B. Zaidel, "Enhancing the cellular downlink capacity via co-processing at the transmitting end," *Proc. IEEE VTC '01 Spring*, vol. 3, pp. 1745–1749, 2001.

[4] T. Weber, I. Maniatis, A. Sklavos, Y. Liu, E. Costa, H. Haas, and E. Schulz, "Joint transmission and detection integrated network (JOINT), a generic proposal for beyond 3G systems," *Proc. ICT'02*, vol. 3, pp. 479–483, 2002.

[5] A. Goldsmith, S. Jafar, N. Jindal, and S. Vishwanath, "Capacity limits of MIMO channels," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 5, pp. 684–702, June 2003.

[6] H. Huang and S. Venkatesan, "Asymptotic downlink capacity of coordinated cellular networks," *Proc. ACSSC '04*, vol. 1, pp. 850–855, 2004.

[7] H. Zhang and H. Dai, "Cochannel interference mitigation and cooperative processing in downlink multicell multiuser MIMO networks," *EURASIP JWCN*, vol. 2004, no. 2, pp. 222–235, 2004.

[8] G. Foschini, H. Huang, K. Karakayali, R. Valenzuela, and S. Venkatesan, "The value of coherent base station coordination," *Proc. CISS '05*, 2005.

[9] M. Karakayali, G. Foschini, and R. Valenzuela, "Network coordination for spectrally efficient communications in cellular systems," *IEEE Wireless Commun. Mag.*, vol. 13, pp. 56–61, 2006.

[10] V. Jungnickel, S. Jaeckel, L. Thiele, L. Jiang, U. Krüger, A. Brylka, and C. Helmolt, "Capacity measurements in a cooperative multicell MIMO network," *IEEE Trans. Veh. Technol.*, vol. 58, pp. 2392–2405, 2009.

[11] TR 36.814 V1.0.0, "Evolved universal terrestrial radio access (e-utra); further advancements for e-utra physical layer aspects," Feb 2009.

[12] W. Zirwas, J. H. Kim, V. Jungnickel, M. Schubert, T. Weber, A. Ahrens, and M. Haardt, *Distributed Antenna Systems*. Auerbach, 2007, ch. 10, Distributed Organization of Cooperative Antenna Systems, pp. 279–311.

[13] Y. Hadisusanto, L. Thiele, and V. Jungnickel, "Distributed base station cooperation via block-diagonalization and dual-decomposition," *Proc. IEEE GLOBECOM' 08*, 2008.

[14] V. Jungnickel, T. Wirth, M. Schellmann, T. Haustein, and W. Zirwas, "Synchronization of cooperative base stations," *Proc. IEEE ISWCS '08*, pp. 329–334, oct 2008.

[15] V. Jungnickel, L. Thiele, M. Schellmann, T. Wirth, W. Zirwas, T. Haustein, and E. Schulz, "Implementation concepts for distributed cooperative transmission," *Proc. ACSSC '08*, oct 2008.

[16] W. Zirwas, W. Mennerich, M. Schubert, L. Thiele, V. Jungnickel, and E. Schulz, *Cooperative Transmission Schemes*. CRC Press, Taylor and Francis Group, 2009.

[17] L. Thiele, T. Wirth, T. Haustein, V. Jungnickel, E. Schulz, , and W. Zirwas, "A unified feedback scheme for distributed interference management in cellular systems: Benefits and challenges for real-time implementation," *Proc. EUSIPCO'09*, 2009.

[18] 3GPP TS 36.300 v8.2.0, "Technical specification group radio access network; evolved universal terrestrial radio access (e-utra) and evolved universal terrestrial radio access network (e-utran); overall description; stage 2," Tech. Rep., 9 2007.

[19] V. Jungnickel, M. Schellmann, A. Forck *et al.*, "Demonstration of virtual MIMO in the uplink," *Proc. IET Smart Ant. and Coop. Commun. Seminar*, 2007.

[20] T. Wirth, V. Jungnickel, A. Forck *et al.*, "Realtime multi-user multi-antenna downlink measurements," *Proc. IEEE WCNC '08*, 2008.

[21] V. Jungnickel, M. Schellmann, L. Thiele, T. Wirth, T. Haustein, O. Koch, E. Zirwas, and E. Schulz, "Interference aware scheduling in the multiuser MIMO-OFDM downlink," *IEEE Commun. Mag.*, vol. 47, pp. 56–66, 2009.

[22] V. Jungnickel, L. Thiele, T. Wirth *et al.*, "Coordinated multipoint trials in the downlink," *Proc. IEEE Globecom Workshops '09*, 2009.

[23] S. Jaeckel, L. Thiele, A. Brylka, L. Jiang, and V. Jungnickel, "Intercell interference measured in urban areas," *Proc. IEEE ICC '09*, 2009.

[24] L. Thiele, T. Wirth, M. Schellmann, Y. Hadisusanto, and V. Jungnickel, "MU-MIMO with localized downlink base station cooperation and downtilted antennas," *Proc. IEEE ICC'09 Workshop on LTE Evolution*, 2009.