

Load-aware reconfiguration of LTE-antennas: dynamic cell-phone network adaptation using organic network control

Sven Tomforde, Alexander Ostrovsky, Jörg Hähner

Angaben zur Veröffentlichung / Publication details:

Tomforde, Sven, Alexander Ostrovsky, and Jörg Hähner. 2014. "Load-aware reconfiguration of LTE-antennas: dynamic cell-phone network adaptation using organic network control." In *11th International Conference on Informatics in Control, Automation and Robotics (ICINCO), September 1-3, 2014, Vienna, Austria*, edited by Joaquim Filipe, Oleg Gusikhin, Kurosh Madani, and Jurek Sasiadek, 236–43. Setúbal: SciTePress.
<https://doi.org/10.5220/0005045102360243>.

Load-aware Reconfiguration of LTE-Antennas

Dynamic Cell-phone Network Adaptation Using Organic Network Control

Sven Tomforde¹, Alexander Ostrovsky² and Jörg Hähner¹

¹Organic Computing Group, University of Augsburg, Eichleitnerstr. 30, 86159 Augsburg, Germany

²Technische Universität München, Arcisstr. 21, 82220 München, Germany

Keywords: Organic Computing, Adaptivity, Intelligent System Control, Learning, Antennas, LTE.

Abstract: The utilisation of cell phone networks increases continuously, especially driven by the introduction of new mobile services and smart phones. Network operators can follow two directions to deal with the problem: either install new hardware or increase the efficiency of the existing infrastructure. This paper presents a novel algorithm to improve the efficiency of current networks by allowing for a self-organised load-dependent reconfiguration of antennas. The algorithm is capable of identifying hotspot traffic, assigning this to a neighbouring cell, and learning the best strategy at runtime. This leads to a self-improving intelligent control mechanism. The simulation-based evaluation results demonstrate the potential benefit, while simultaneously keeping the hardware's deterioration at a comparable level.

1 INTRODUCTION

Wireless cellular networks are growing rapidly. Cisco estimates that the overall mobile traffic in 2017 will reach 11.2 exabytes per month, which is 13 times more than it was in the year 2012¹. As a result, network operators have to increase the capacity of their networks significantly. Since new hardware installations are costly, intelligent control mechanisms and means to optimise the utilisation of existing infrastructure are necessary. Such an approach is investigated by this paper.

Typically, the load within a cell phone network such as LTE (Long Term Evolution) is not homogeneously distributed – instead, it is subject to spatial and temporal variations (Willkomm et al., 2009). While some cells are overloaded at one point of time, they can be lightly loaded at some other point of time. For instance, this can be observed in office areas, where heavy traffic load appears at working hours followed by light usage at other times. Such an uneven loading can also be observed among neighbouring cells at the same point in time. The approach presented in this paper provides a solution that balances the load among neighbouring LTE cells without the need of major hardware modifications of the antennas

or changes in the LTE specifications. Thereby, the algorithm is able to improve its behaviour over time based on an online learning approach.

The remainder of this paper is concerned with the developed algorithm. Therefore, we start with an overview of the current state of the art in Section 2, followed by a description of the physical model used as basis for this paper (Section 3). Section 4 introduces the novel algorithm for dynamic antenna reconfiguration. Afterwards, its performance is evaluated based on simulations (Section 5). Finally, the paper closes with a summary and an outlook in Section 6.

2 STATE OF THE ART

Methods for optimising antenna parameters for UMTS (Universal Mobile Telecommunications System), LTE and other mobile networks have been widely discussed in the literature. In (Temesvary, 2009), an algorithm for the optimisation of antenna tilt and power based on LTE networks is presented, which makes use of the optimisation heuristic *Simulated Annealing*. The goal is to improve SINR (Signal to Interference plus Noise Ratio) measured at the UEs (User Equipment; i.e. cell phones). Therefore, so-called CQI (Channel Quality Indication) reports are

¹Cf. “Rethink Research” (CISCO), http://www.theregister.co.uk/2013/02/11/mobile_traffic_will_be_video/, 2013

collected². The simulation-based evaluation showed that employing a combination of antenna power and tilt optimisation does not lead to significantly better performance than just optimising the tilt.

Furthermore, the work in (Deruyck et al., 2013) demonstrated that decreasing transceiver power results not necessarily in a significant decrease of the antenna's overall power consumption. Another attempt to optimise the cell network's efficiency has been presented in (Du et al., 2002): A genetic algorithm has been used to determine size and shape of cells. Thereby, antenna gains are optimised in each direction to find a trade-off between minimising the overall base station power consumption and maximising the capacity. Besides cell shapes, so-called *cell-zooming* has been investigated in (Niu et al., 2010). The concept relies on switching inactive cells off, which results in saving energy due to a concentration among only few necessary cells. Simulations carried out here show that about 30 to 50% of the base stations can be switched off without loss of functionality – but a transfer to UMTS or LTE networks is missing.

More focussed towards the algorithm presented in this paper, (Awada et al., 2011) investigates the usage of Taguchi's Method (Weng et al., 2007) for the optimisation of uplink power, antenna tilt and azimuth. Simulation results showed that an offline optimisation converges faster than approaches using simulated annealing (Kirkpatrick et al., 1983) in most cases. Similarly, (Razavi, 2012) focuses on the antenna tilt as optimisation parameter by using the golden section search algorithm to find an optimised angle, followed by frequent explorations to fine-tune it. The results for a homogeneous traffic distribution show that the optimal antenna tilt is rather large, so this method converges fast. In (Kim et al., 2012), the authors model a mobile network as a M/G/1 queue and introduce a distributed algorithm to optimise parameters such as the throughput of the network. The algorithm is shown to converge fast towards the searched optimum but has not been applied to UMTS or LTE networks, yet. Similarly, the authors in (Fehske et al., 2013) model a LTE network as a M/M/1 queue and introduce a centralised algorithm to optimise handover parameters and antenna tilts. A system-level simulation shows that it is able to improve user throughput even during low-traffic times. In further work, e.g. (Razavi et al., 2010), reinforcement learning techniques are applied to improve coverage and capacity aspects.

The approach presented in this paper is different

²An overview of metrics related to CQI can be found in a Technical Report by Ericsson, available online: <http://www.ericsson.com/res/docs/whitepapers/wp-lte-acceptance.pdf>

to the afore mentioned work due to several reasons. The purpose is to automatically relieve hotspot traffic in overloaded cells during runtime (i.e. while the antennas operate), while most of the existing work is situated at design-time. A *hotspot* is an accumulation of UEs that lasts for a certain period of time. Relieving this traffic is done by shifting it to a neighbouring cell with less load. In contrast to our solution, existing approaches for this problem require many steps until a good configuration is found. Tilting antennas has impact on the hardware: Doing this too often will lead to increased maintenance intervals and necessary exchange of components.

Finally, the presented work is part of the *Organic Network Control* project (ONC) (Tomforde et al., 2009). Based on principles of Organic Computing (Müller-Schloer, 2004), the project investigates possibilities to augment data communication networks with “life-like” characteristics, i.e. self-organisation, robustness, and flexibility. The first phase was concerned with self-improving reconfiguration of existing network protocol parameters in response to changing conditions (Tomforde and Hähner, 2011). The current second phase shifts the focus towards reconfiguration of hardware and collaborative solutions.

3 ANTENNA TILT AND PHYSICAL MODEL

The term *antenna tilt* describes the angle between the antenna's main beam and the horizontal pane. When the beam is directed downwards, the antenna is tilted down; when the beam is directed upwards, the antenna is tilted up. By convention, a negative angle indicates tilting the antenna up and a positive tilting it down – a angle of 0 means that the beam is parallel to the horizontal pane (Bratu, 2012). The adjustment of tilts can be achieved either mechanically, electrically (Bratu, 2012), or by vertical beam forming (Nokia Siemens Networks Corporation, 2012). Mechanical tilts are a result of mounting antennas with a certain angle. In contrast, the electrical tilt is adjusted by changing the phase characteristics – this can be done remotely using *Remote Electrical Tilt* (RET). *Vertical Beam Forming* adjusts the tilts for multiple UEs independently (has to be supported by the antenna's hardware and specification). Tilt adjustment provides several different possibilities for optimisation. However, changing antenna parameters is not trivial and wrong decisions may lead to interferences and a decreased coverage (Holma and Toskala, 2012).

The *Physical Model* is used to predict the propagation of radio waves. The distribution of these ra-

radio waves are influenced e.g. by obstacles and the atmosphere. When a radio wave impinges an object, it can pass through it, be absorbed, or it can be reflected, scattered (i.e. reflection to multiple directions) or diffracted (Dean, 2009). Wireless signals can follow multiple paths (*multipath* characteristic) – it is therefore difficult to predict the exact behaviour. However, this can be approximated by combining pathloss, shadow fading, and fast fading (Ghosh et al., 2010). *Pathloss* means the damping that occurs in relation to the distance passed by the signal and can be approximated (in dB) for macro cells in urban area as follows (3GPP, 2012b):

$$L(R) = 40 \times (1 - 4 \times 10^{-3} \times D_{hb}) \times \log_{10}(R) - 18 \times \log_{10}(D_{hb}) + 21 \times \log_{10}(f) + 80 \quad (1)$$

where R is the distance between the base station and the UE (in km), f is the carrier frequency (in MHz), and D_{hb} is the height of the base station above average rooftop level (in m).

The pathloss model described above assumes that the damping is constant for all paths. This assumption does not hold for all cases: While some paths suffer increased loss (e.g. due to buildings), others are less obstructed. This effect is called *shadow fading* (Dean, 2009) and can be critical on cell edges and create coverage holes. Models for shadow fading use a log-normal distribution (Ikuno et al., 2010). Hence, the combined effect (L in dB) of pathloss and shadow fading can be expressed as: $L = \bar{L} + X$, where \bar{L} is the mean pathloss, and X is a normal distributed random variable with a mean of 0 and a standard deviation of 10 (3GPP, 2012b; Ikuno et al., 2010). Due to changes in the topology and vegetation, shadow fading changes over time (Wang, 2007). Contrary to intuition, rain, fog and snow have only a negligible effect on signal damping (Wang, 2007).

Antenna tilt and azimuth (i.e. the angle between the antenna's main beam and the vertical plane) have also impact on the signal damping. Decreasing the vertical angle between UE and the eNodeB (the particular E-UTRAN Node with the considered antenna) – in comparison to the angle with maximum gain direction – will also lead to a decrease in the signal damping. The gain of antenna power in a given direction is contrary to an antenna that radiates equally in all directions (isotropic radiator) (Hill, 1976). Taking this into account, the received power can be estimated as follows (3GPP, 2012a):

$$RX_{PWR} = TX_{PWR} - \max(L - G_{TX} - G_{RX}, MCL) \quad (2)$$

where RX_{PWR} is the received power, TX_{PWR} the transmitted power, L the pathloss, G_{TX} the transmitter antenna's gain, G_{RX} the receiver antenna's gain, and

MCL the minimum coupling loss (which is defined as 70dB for urban areas). Temporary anomalies that may disturb the radio wave propagation (i.e. tropospheric ducting) are neglected in the context of this paper. The algorithm presented in the following considers this physical model.

4 DYNAMIC ANTENNA RECONFIGURATION

This section describes the distributed algorithm for the optimisation of congested cells. It reconfigures antenna tilts such that possible hotspots are shifted from the coverage area of the congested cell to the coverage area of a neighbouring (underutilised) cell. Down-tilting should lead to a decrease in the covered area and vice-versa – due to physical and weather conditions, this is not always the case. Therefore, the algorithm is based on estimating the achieved success. This is combined with reinforcement learning concept to improve this behaviour at runtime.

The basic idea of the algorithm is to deal with the existing hardware and to operate without changes in the LTE specifications. Important mechanisms are already available: 1) *antenna tilts* can be changed with *Remote Electrical Tilt* (RET), 2) the *discovery of neighbours* can be done with *Automated Neighbour Relation* (ANR) (3GPP, 2012b), 3) the *communication between neighbours* is implemented using the X2 interface (3GPP, 2012a), and the *positioning of users* is supported by LTE (Iwamura et al., 2009).

The algorithm for online antenna tilt optimisation consists of five parts: the basic algorithm is responsible for optimising the mapping of UEs to eNodeBs (Part 1: Optimisation). This requires further aspects: the identification of hotspot traffic that fulfils the requirements to be handled as a cluster by the algorithm (Part 2: Identification), a method to select a neighbouring cell to relieve the cluster to (Part 3: Neighbour Policy), a mechanism to learn from previous experiences (Part 4: Learning), and finally a measure to quantify the similarity of two clusters (Part 5: Similarity). The remainder of this section introduces these five parts in detail and discusses the possibilities and limitations of the approach.

4.1 Part 1: Optimisation

The optimisation part of the algorithm aims at relieving clusters. Therefore, it analyses data collected in previous runs and data provided by the exploration part (see Section 4.3). This analysis leads to a prediction of which neighbour should be tilted up. If no

such data exists for the analysis, the algorithm uses heuristics to generate the prediction. In order to keep the configuration chosen by the operator as static as possible (i.e. apply as few changes simultaneously as possible), only one neighbour is taken into consideration in each step. The following Algorithm 1 is executed for each eNodeB.

In Algorithm 1, the variable $Cell_i$ is the particular cell maintained by the antenna. The parameter Deg defines the number of degrees an antenna is tilted and can be adjusted by the operator – it should be small to prevent large deviations from the initial configuration. The policy $Pi_{Cluster}$ defines which cluster is to be processed – a simple policy is to select a cluster that is located as far as possible from the serving eNodeB. Thereby, CL is the currently investigated cluster. $CandNeighbour$ specifies the set of possible neighbouring cells to apply hotspot traffic to and N_s is the currently investigated neighbour out of this set. Finally, Fit_i estimates the performance (or *fitness*) of the cell before tilting, Fit_s estimates the performance after the last tilt change has been applied.

Algorithm 1: Dynamic Antenna Reconfiguration.

```

1: Input:  $Deg, Pi_{Cluster}, Pi_{Neighbour}$ ;
2: if  $Cell_i$  is congested and hotspots exist then
3:    $CL \leftarrow$  Select Clusters according to policy  $Pi_{Cluster}$ ;
4:    $Fit_i \leftarrow$  Current performance of cell;
5:   for all  $CL_i \in CL$  do
6:      $CandNeighbours \leftarrow$  All Neighbours with free capacity near  $CL_i$ ;
7:      $N_s \leftarrow$  Select eNodeB from  $CandNeighbours$  according to policy  $Pi_{Neighbour}$ ;
8:     Tilt down eNodeB serving Cluster by  $Deg$  degrees;
9:     Tilt up  $N_s$  by  $Deg$  degrees;
10:     $Fit_s \leftarrow$  Save to what extend Cluster was relieved;
11:    if  $Fit_i \geq Fit_s$  then
12:      Reset tilts of eNodeB and  $N_s$ ;
13:    end if
14:  end for
15: end if

```

Only neighbouring antennas with free capacity to serve the cluster should be considered for tilting up (otherwise they may also overload). Furthermore, a relieved cluster does not necessarily result in a better performance at user side. Hence, the tilts are reset if the end user performance decreases.

4.2 Part 2: Identification of Clusters

The algorithm as presented before relies on knowledge about existing clusters of UEs. Hereby, the number of clusters is unknown. Different approaches to identify an unspecified number of clusters are known in literature, DBSCAN (Density-Based Spatial Clustering of Applications with Noise – (Ester et al.,

1996)) and OPTICS (Ordering Points To Identify the Clustering Structure – (Ankerst et al., 1999)) are the most appropriate ones in the context of this paper. Both approaches identify clusters based on the density of points. Therefore, they need three parameters: 1) *SetOfPoints* (data for clustering – here: positions of UEs), 2) *Eps* (maximum distance of the points within a single cluster), and 3) *MinPts* (minimum number of points to form a cluster). These three parameters must be known in advance; a good setup is chosen empirically. In contrast to DBSCAN, OPTICS does not return a certain cluster, but an ordering of possible cluster candidates. Hence, we used OPTICS as techniques to identify cluster.

4.3 Part 3: Selection of Neighbours

Within Algorithm 1, a policy $Pi_{Neighbour}$ is needed that chooses a neighbouring antenna for tilting up (and for taking over the hotspot traffic). Depending on previous experiences, this policy has different options to pursue. These experiences are either provided by the exploration algorithm or exist due to previous runs of the optimisation algorithm. In case there is no previous cluster information being similar to the current cluster CL_i , no configuration can be re-used. A conservative approach is here to keep the current tilt settings until the exploration phase provides results. In contrast, an opportunistic approach is to choose a neighbour eNodeB heuristically. Both approaches have advantages and drawbacks. The eNodeB chosen by a heuristic may be a good choice, but its tilt change may deteriorate the channel quality of the users without relieving the cluster. If knowledge with clusters similar to CL_i has been collected before, the policy $Pi_{Neighbour}$ determines the tilt configuration that relieves CL_i according to the previous experiences.

Two different heuristics have been implemented and tested in simulations: 1) select a random neighbour which azimuth is towards the cluster, and 2) choose the eNodeB which has an azimuth towards CL_i and a part of the cluster in its coverage area (this is based on the observation that a cluster located at the edge of a cell is often located within the coverage areas of two cells). A hybrid solution consisting of both is to use 1) if 2) has not been successful.

4.4 Part 4: Learning

In the context of the optimisation algorithm, learning is concerned with the policy $Pi_{neighbour}$. The selection of an appropriate neighbouring antenna to hand over hotspot traffic should be improved over time by considering the success of previous actions. Until now,

the algorithm selects a neighbour randomly or according to heuristics. Afterwards, information whether the tilt change led to a relief of a cluster or not (based on the performance estimation) is collected. The more UEs switched from the cluster to a non-congested cell, the better is the quality of the tilt change. This can be formalised as follows:

$$q = |U_b| - |U_a| \quad (3)$$

with q being the performance function, U_a the set of UEs served by the particular overloaded antenna before the tilt changes have been applied, $|U_a|$ the number of UEs contained in this set, and U_b the set of UEs served by the particular overloaded antenna after the tilt change. While running, the algorithm collects the qualities of tilt changes and stores them as quadruples containing cluster information, neighbour to cooperate with, quality of the tilt change and time. To decide which neighbour should be tilted up for a given cluster, the qualities of tilt changes for all UEs are combined using the beta distribution density function. Especially in the context of deriving reputation values for participants in e-commerce systems, the beta distribution density function is used due to its simplicity as well as solid mathematical foundations (Josang and Ismail, 2002).

In the context of this paper, the beta distribution approach is used to predict to which extend tilting up of an antenna will relieve a cluster. Therefore, previous data about tilt changes has to exist. Assume we want to determine whether tilting up of a neighbour N_j will relieve a cluster CL_i . We already collected historical data to what extend tilting up of N_j relieved a previously observed similar cluster CL_i . First, we set the two parameters as needed for the distribution function α and β to a predefined constant value larger than one. We know due to the properties of the beta distribution PDF that the modus is at 0.5. So we assume 0.5 to be a neutral value. Afterwards, we iterate through the historical data. Each time an indication occurs that a tilt change of N_j relieved CL_i , we increment α . In contrary, each time an indication occurs that a tilt change of N_j did not relieve CL_i , we increment β . Due to the properties of the beta distribution PDF, we know: When α is larger than β , the modus increases (and vice-versa). Hence, if there are more positive indications than negative ones, the modus is larger than 0.5. For further improvements, the increments can also be weighted by the quality of the particular tilt changes. The larger the absolute difference is, i.e. $|\alpha - \beta|$, the larger is the absolute value of the difference between the modus and 0.5. Hence, we can estimate the probability p that tilting up of N_j will relieve CL_i by calculating the absolute difference of the x-coordinate of the maximum and 0.5.

We observe that the bell shaped curve becomes broader with decreasing values of α and β and narrower with increasing values of α and β . Hence, the value of the p %-quantile is used with a small value for p instead of the modus. If p is small enough, then the p %-quantile is smaller than the modus – if the modus is larger than or equal to 0.5. With larger values for α and β , the distance between p %-quantile and the modus decreases.

By omitting the normalisation factor of the original beta distribution function, the computability can be improved. The result is given in the following function:

$$f(p, \alpha, \beta) = P^{\alpha-1} \times (1-p)^{\beta-1} \quad (4)$$

for an interval $[0; 1]$ with $0 \leq p \leq 1$, $\alpha > 0$ and $\beta > 0$. Thereby, f is the distribution function, α and β the weighting factors and p a constant.

4.5 Part 5: Similarity of Clusters

The approach as presented before relies on previous experiences with similar clusters of hotspot traffic. This implies the possibility to compare clusters and to store information about clusters. In an ideal case, two clusters will consist of users at exactly the same positions at different points of time. In reality, this will not happen.

Assume we have two clusters c and c_0 occurring at two different points of time. When the positions of the UEs in c and c_0 only slightly differ, we define c and c_0 as similar. Then e.g. experiences with c_0 can be used to predict an eNodeB for tilting up to relieve c . To compare two clusters we use a modified Fowlkes-Mallows index (Fowlker and Mallows, 1983):

$$FM = \sqrt{\frac{P_{c,c'}}{P_{c,c'} + P_c} \times \frac{P_{c,c'}}{P_{c,c'} + P_{c_0}}} \quad (5)$$

with $P_{c,c'}$ being the UEs that are contained in both sets (i.e. UEs that are contained in the convex hull of the cluster), P_c the positions of UEs contained only in c , and P_{c_0} the positions of UEs contained only in c_0 .

5 EVALUATION

5.1 Experimental Setup

To test the implemented algorithm, simulations were performed using a modified *Vienna LTE System Level Simulator* (Ikuno et al., 2010). The simulator was configured using the values as listed in Table 1.

Table 1: Simulation parameters.

Antenna Model	Kathrein 742 215
Base station height	20 m
Transmit power	30 dB
Mobile height	1.5 m
Inter eNodeB distance	500 m
Pathloss model	3GPP TR 36.942 (see (3GPP, 2012b))
Shadow fading model	Lognormal distributed, 2D space correlated (Claussen, 2005)
Electrical tilt range	0° – 10°
Mechanical tilt	0°
Channel Model	Winner II+
Min. coupling loss	70 dB
Scheduler	Round Robin
MIMO	2 senders, 2 receivers
MIMO transmission	Closed-loop spatial MUX
Bandwidth	20 Mhz / 100 RBs
Traffic distribution	FTP: 10%; HTTP: 20%; VIDEO: 20%; VoIP: 30%; Gaming: 20% (3GPP, 2007)
Hotspot definition	≥ 3 UEs in 3 m
Min. switching UEs	5 UEs

5.2 Experimental Results

A: Antenna Tilt Reconfiguration

The developed algorithm is tested on different scenarios. Each scenario consists of 21 active eNodeBs and 36 passive eNodeBs. The scenarios differ in terms of UEs served by each eNodeB. Furthermore, hotspot traffic is simulated – which has to be relieved by the algorithm. For this paper, we investigated three different scenarios – where each eNodeB serves up to 9 (Scenario 1 – results are listed in Table 2), up to 15 (Scenario 2 – results are listed in Table 2), and up to 30 UEs (Scenario 3 – results are listed in Table 3). Each randomly generated scenario is initially generated without hotspots. Afterwards, this is optimised using simulated annealing with the goal function to increase the 10%-percentile throughput. Thereby, “no improvements for 20 iterations” has been chosen as termination criterion. This optimised setting then serves as input for the algorithm, limiting the optimisation potential to just the additional hotspot traffic. Therefore, two random hotspots with 10 to 40 UEs and a radius between 20m and 70m are added. The distance between the centre of a hotspot and the nearest eNodeB is at least 175m – in case the two hotspots are located too close to each other, they automatically merge into one large hotspot. In contrary, a hotspot can automatically split into parts if the distance between UEs is too large. Similarly, a hotspot which is too sparse is not recognised by the clustering algorithm (cluster: min. 10 UEs; at least 25m to the nearest neighbour).

The clustering part of the algorithm is configured

to only accept clusters defined by the area of the convex hull of up to 300m². Thereby, a single hotspot can be relieved multiple times by different antenna configurations. The results as described in the following reflect aggregated simulation results of several (in most cases: 10) randomised scenarios. In all cases, we distinguish between two possibilities: Either we let the algorithm modify the tilt by an angle of 1° or by an angle of 2°. We measured the improvement of the 10%-percentile throughput of the overloaded eNodeB and its neighbours as an aggregated value, since an improvement for the eNodeB might result in a worsening for the neighbours accepting the hotspot traffic (and vice-versa). The results given in the tables reflect averages over all runs of the simulation.

Table 2 lists the results of the evaluation for the first scenario with six to nine UEs connected to each eNodeB additionally to the hotspots. Thereby, the abbreviation RB stands for *Resource Block*. When setting the angle of tilt change to 1°, the mean improvement of the 10%-percentile throughput is significantly larger compared to the results with a tilt change of 2°. This is probably due to the fact that there is a higher negative effect on the channel quality for a higher change of the tilts.

The same simulations have been evaluated for the scenario with 15 UEs per eNodeB (in addition to the hotspot traffic to be relieved). Table 2 lists the achieved results. Again, changing the antenna tilt by 1° led to a better performance than the change by 2°. Compared to the previous results, we can observe that the algorithm needed slightly more iterations (2.13 vs. 1.74 for 1° change; 2.22 vs. 1.81 for 2° change). This is still extremely fast – especially compared to the usage of an optimisation heuristic such as simulated annealing, where hundreds of iterations are needed. The simulation results for 30 UEs per eNode support the observations for the previous two scenarios, see Table 3 for details.

In some scenarios shifting the hotspot did not lead to an increased 10%-percentile throughput for the overloaded eNodeB and its neighbours (cf. the number of successful relieves in Table 2 and 3). There are several reasons for this behaviour. First, shifting a hotspot may overload the neighbour. This can happen when the number of users in the hotspot is large and the hotspot is near the coverage area of the neighbour that should tilt up. A second reason is that UEs might move to another eNodeB than the selected partner for the handover process. This can especially happen in case of large clusters. Finally, UEs connected to the overloaded eNodeB that are not in the cluster may switch to other neighbouring cells after changing tilts – which again influences the evaluation results.

Table 2: Simulation results for the optimisation algorithm for 6 and 15 UEs per eNodeB additionally to the clusters.

Parameter	Tilt change 1° [6 UEs]	Tilt change 2° [6 UEs]	Tilt change 1° [15 UEs]	Tilt change 2° [15 UEs]
Number of hotspots	276	276	218	218
Number of scenarios	241	241	177	177
Max. number of RB scheduled to each UE in overloaded cell	3	3	3	3
Min. number of RB scheduled to each UE in free cell	6	6	6	6
Mean improvement of 10%-percentile throughput	17.38%	7.62%	5.04%	1.06%
Range of improvements (min. and max.)	[-55.61%; 186.08%]	[-67.74%; 174.91%]	[-37.54%; 79.67%]	[-49.89%; 119.57%]
Mean CQI decrease	0.27	0.34	0.37	0.40
Mean number of iterations until convergence	1.74	1.81	2.13	2.22
Mean distance of cluster centre to eNodeB	247.79 m	241.57 m	245.01 m	243.27 m
Number of successful relieves	198	400	111	187

Table 3: Simulation results for the optimisation algorithm for 30 UEs per eNodeB additionally to the clusters.

Parameter	Tilt change 1°	Tilt change 2°
Number of hotspots	276	276
Number of scenarios	225	225
Max. number of RB scheduled to each UE in overloaded cell	2	2
Min. number of RB scheduled to each UE in free cell	3	3
Mean improvement of 10%-percentile throughput	8.71%	-2.63%
Range of improvements (min. and max.)	[-34.58%; 69.81%]	[-46.45%; 53.94%]
Mean CQI decrease	0.20	0.31
Mean number of iterations until convergence	1.25	1.55
Mean distance of cluster centre to eNodeB	257.64 m	247.65 m
Number of successful relieves	121	294

B: Number of Users, CQI and Throughput

The last part of the evaluation deals with the impact of the channel quality on the throughput. When increasing the channel quality, the throughput usually increases as well. This is illustrated by Figure 1 (see the red trend line). Here, CQI before and after the optimisation is compared to the throughput before and after the optimisation – if the hotspot was successfully relieved.




Figure 1: CQI and 10%-percentile throughput.

The correlation between the number of users switched to the neighbour and the increase of the throughput is less obvious, see Figure 2. The results show that the 10%-percentile throughput increases, if there are at least 5 UEs switching from an overloaded cell to a free neighbour. An advantage of tilt-based load balancing (i.e. in comparison to only handover-based load balancing) is the possible limitation of the channel quality's decrease (or even the increase for UEs shifted to another cell).

6 CONCLUSION

This paper presented a novel distributed algorithm for




Figure 2: Number of users switched from the cluster in the overloaded cell to a free neighbour vs. increase of 10%-percentile throughput.

the optimisation of antenna tilts in LTE networks. Thereby, hotspots of users are shifted from an overloaded cell to a free neighbouring cell, which requires the LTE protocol in release 10 and higher (LTE-Advanced). The simulation-based evaluation demonstrated the potential benefit of assigning clusters of users to neighbouring cells in oversaturated conditions. Thereby, the search for an optimal antenna tilt configuration is not trivial due to the unpredictable propagation of radio waves – which leads to the demand of an intelligent control mechanism. Such a control mechanism is presented by this paper that combines the advantages of a self-organised approach and the capability of learning at runtime.

While the approach presented before optimises hotspot traffic, future work will focus on a generalisation towards relieving different kinds of congestions. Therefore, cells can be split into geographical sectors and the tilt change can be explored for each sector independently. Based on this method, users may be clustered not only geographically but also based on their channel quality. However, when changing antenna tilt, the SINR of each cluster does not change

homogeneously.

Furthermore, the presented algorithm does not consider the demanded QCI and throughputs of different users. Real user data is needed to verify whether there are clusters of users who need a high throughput (e.g. privileged users) and clusters of users who do not need a high throughput. In combination with handover parameters, the scheduling algorithm, MIMO transmission techniques and other parameters, the optimisation might be even more successful.

REFERENCES

- 3GPP (2007). LTE Physical Layer Framework for Performance Verification. Tech. rep., 3rd Generation Partnership Proj.
- 3GPP (2012a). Evolved Universal Terrestrial Radio Access Network; X2 general aspects and principles. Tech. rep. v. 11.0.0, 3rd Generation Partnership Proj.
- 3GPP (2012b). Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Radio Access; Radio Frequency system scenarios. Tech. rep. v. 11.0.0, 3rd Generation Partnership Proj.
- Ankerst, M., Breunig, M., Kriegel, H., and Sander, J. (1999). Optics: Ordering points to identify the clustering structure. In *Proc. of SIGMOD99*, pages 49 – 60.
- Awada, A., Wegmann, B., Viering, I., and Klein, A. (2011). Optimizing the Radio Network Parameters of the LTE System Using Taguchi's Method. *IEEE Trans. on Vehicular Technology*, 60(8):3825–3839.
- Bratu, V.-I. (2012). Self-optimisation of Antenna Tilt in Mobile Networks. Master's thesis, KTH Royal Institute of Technology.
- Claussen, H. (2005). Efficient Modelling of Channel Maps with Correlated Shadow Fading in Mobile Radio Systems. In *Proc. of 16th Symp. on Personal, Indoor and Mobile Radio Com.*, pages 512 – 516.
- Dean, T. (2009). *Network+ Guide to Networks*. Cengage Course Technology.
- Deruyck, M., Joseph, W., Lannoo, B., Colle, D., and Martens, L. (2013). Designing Energy-Efficient Wireless Access Networks: LTE and LTE-Advanced. *IEEE Internet Computing*, 17(5):39 – 45.
- Du, L., Bigham, J., Cuthbert, L., Parini, C., and Nahi, P. (2002). Cell size and shape adjustment depending on call traffic distribution. In *Proc. of WCNC'02*, volume 2, pages 886 – 891.
- Ester, M., Kriegel, H. P., Sander, J., and Xu, X. (1996). A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *Proc. of Int. Conf. on Knowledge Discovery and Data Mining*, pages 226 – 231.
- Fehske, A., Klessig, H., Voigt, J., and Fettweis, G. (2013). Concurrent Load-aware Adjustment of User Association and Antenna Tilts in Self-organising Radio Networks. *IEEE Trans. on Vehicular Technology*, 1:99 ff.
- Fowlker, E. and Mallows, C. (1983). A Method for Comparing Two Hierarchical Clusterings. *J. of the American Statistical Association*, 78(383):553 – 569.
- Ghosh, A., Zhang, J., Andrews, J., and Muhamed, R. (2010). *Fundamentals of LTE*. Prentice Hall.
- Hill, J. (1976). Gain of directional antennas. Technical report, Watkins-Johnson Co.
- Holma, H. and Toskala, A. (2012). *LTE Advanced: 3GPP Solution for IMT-Advanced*. Wiley and Sons.
- Ikuno, J. C., Wrulich, M., and Rupp, M. (2010). System Level Simulation of LTE Networks. In *Proc. of VTC'10*, pages 1 – 5.
- Iwamura, M., Umesh, A., and Hapsari, W. A. (2009). Further Enhancements of LTE – LTE Release 9. *NTT Docomo Tech. J.*, 12(1):45 – 53.
- Josang, A. and Ismail, R. (2002). The Beta Reputation System. In *Proc. of Bled Electronic Commerce Conf.*, pages 41 – 55.
- Kim, H., de Veciana, G., Yang, X., and Venkatachalam, M. (2012). Distributed Alpha-optimal User Association and Cell Load Balancing in Wireless Networks. *IEEE/ACM Trans. on Networks*, 20(1):177–190.
- Kirkpatrick, S., Jr., D. G., and Vecchi, M. P. (1983). Optimization by Simulated Annealing. *Science*, 220(4598):671 – 680.
- Müller-Schloer, C. (2004). Organic Computing: On the Feasibility of Controlled Emergence. In *Proc. of CODES and ISSS'04*, pages 2–5.
- Niu, Z., Wu, Y., Gong, J., and Yang, Z. (2010). Cell zooming for cost-efficient green cellular networks. *IEEE Com. Mag.*, 48(11):74–79.
- Nokia Siemens Networks Corporation (2012). Active antenna systems: A step-change in base station site performance. Technical report.
- Razavi, R. (2012). Self-optimisation of antenna beam tilting in lte networks. In *Proc. of VTC'12*, pages 1–5.
- Razavi, R., Klein, S., and Claussen, H. (2010). Self-optimisation of Capacity and Coverage in LTE Using a Fuzzy Reinforcement Learning Approach. In *Proc. of IEEE PIMRC'10*, pages 1865 – 1870.
- Temesvary, A. (2009). Self-Configuration of Antenna Tilt and Power for Plug and Play Deployed Cellular Networks. In *Proc. of Wireless Communications and Networking Conf.*, pages 1–6.
- Tomforde, S., Cakar, E., and Hähner, J. (2009). Dynamic Control of Network Protocols - A new vision for future self-organised networks. In *Proc. of ICINCO'09*, pages 285 – 290.
- Tomforde, S. and Hähner, J. (2011). *Biologically Inspired Networking and Sensing: Algorithms and Architectures*, chapter Organic Network Control – Turning standard protocols into evolving systems, pages 11 – 35. IGI Publishers.
- Wang, J. (2007). *Performance Analysis and Measurement of CDMA2000 System in Metropolitan Taipei*. PhD thesis, Nat. Taiwan Univ. of Science and Tech.
- Weng, W., Yang, F., and Elsherbeni, A. (2007). Linear Antenna Array Synthesis Using Taguchi's Method: A Novel Optimization Technique in Electromagnetics. *Trans. on Antennas and Propagation*, 55:723 – 730.
- Willkomm, D., Machiraju, S., Bolot, J., and Wolisz, A. (2009). Primary User Behavior in Cellular Networks and Implications for Dynamic Spectrum Access. *IEEE Com. Mag.*, 47(3):88 – 95.