

A Two-Stage Approach to WLAN Planning: Detailed Performance Evaluation Along the Pareto Frontier

Andreas Eisenblätter¹ Hans-Florian Geerdes¹ James Gross² Oscar Puñal² Jonas Schweiger¹

¹Zuse Institute Berlin, Optimization Department, {eisenblaetter|geerdes|schweiger}@zib.de

²Mobile Network Performance Group, UMIC Research Centre, RWTH Aachen University, {gross|punal}@umic.rwth-aachen.de

Abstract—This paper proposes an efficient solution to the open problem of network planning for large-scale WLAN deployments. WLAN performance is governed by the CSMA-CA protocol, whose dynamic effects are difficult to capture. Accurate performance evaluation depends on simulations and takes time. A detailed analysis of dozens candidate designs with varying AP positions and channel assignments during network planning is therefore infeasible. In our solution, we first identify few good candidate designs using a multi-criteria optimization model, which features notions of cell overlap and station throughput. These candidate designs are taken from the corresponding Pareto frontier. In the second step, we evaluate the performance of the candidate designs by means of simulations. We apply our method to a realistic, large-scale planning scenario for an indoor office environment. The detailed simulations reveal important characteristics of the candidate designs that are not captured by the optimization model. The resulting performance differs significantly across the candidate designs. Hence, this approach successfully combines the benefits of mathematical optimization and simulations, yet avoiding their individual drawbacks.

I. INTRODUCTION

Wireless Local Area Networks (WLANs) implementing the IEEE 802.11 standard are among the most popular wireless broadband networks today. Apart from the unstructured deployment of single WLAN access points (APs), WLANs consisting of several APs that belong to one administrative domain are deployed increasingly to provide networking services over a larger area. In these cases *Network planning* is one means to optimize performance. The goal is to find an optimal *network design* comprising the placement and configuration (e.g., channel assignment) of the APs. The performance of a WLAN with multiple APs varies largely with their placement and configuration.

A. Challenges

There are several subtleties related to WLAN performance. Capturing those in a closed-form mathematical model that applies to generic network designs is an open problem. Consider, for example, the average goodput per station across the whole

network as performance metric. This quantifies the capability of a WLAN to deliver payload in the presence of interference (among other factors). Interference arises from uncoordinated medium access of APs using the same or overlapping channels. Weak interference (below the carrier sensing range) merely increases the packet error rates of the links and thus degrades performance on the physical layer. Strong interference, in contrast, contributes mainly to the contention on the medium access layer due to the carrier sense multiple access protocol with collision avoidance (CSMA-CA). Taking into account the nonlinearity of goodput models for WLAN contention, it is clear that there is a lack of automatic methods for finding AP locations and channels such that coverage and capacity of large-scale networks can be jointly optimized. Notice the two dimensions of “large-scale.” There is, of course, the size of the planning task measured in the number of APs and their candidate locations. On the other hand, there is the degree to which neighboring cells can interfere with each other. If three dimensional scenarios are considered, there is also the point of cross-floor interference in multi-floor indoor deployments. This grows with the number of floors and adds to the complexity of the problem when compared to two-dimensional outdoor macro-cellular 2G/3G networks, for example. Within multi-floor environments, modeling interference becomes pivotal.

B. Related Work

General methods and modeling techniques for radio network design and optimization (also treating the problems studied here) can be found in [1], [2]. [2] describes an integrated model for channel assignment and base station positioning. However, the specialties of WLAN technology are not addressed.

The articles [3], [4] deal with maximizing coverage while [5], [6] deal with channel assignment. Common to these contributions is that they consider either coverage improvement or interference reduction. Except for [6], the obtained optimization results are hardly analyzed in a more realistic setting than the optimization model itself.

Several contributions consider AP location and channel assignment together. In [7], greedy strategies are proposed for first finding AP locations and then assigning channels. In [8], coverage maximization is complemented by simultaneously

This work is partially supported by the DFG Excellence-Cluster “UMIC – Ultra-high speed mobile information and communication” and by the DFG research center MATHEON “Mathematics for key technologies.”

checking for a valid channel assignment in a mathematical programming model. A detailed, non-linear optimization model taking AP placement and channel assignment into account is proposed in [9]. The authors propose an explicit exhaustive search for finding good designs. In [10], the authors study the applicability of genetic algorithms to the problem. Their model is limited to two-dimensional planning problems.

An approach based on purely minimizing SINR by AP placement and channel assignment is studied in [11]. Finally, a mathematical model for integrated planning is proposed in [12]. None of the above works provides an analysis of their findings with respect to a realistic WLAN performance model—neither by simulations nor by experiments.

C. Contributions and Organization of the Paper

This article contributes in two ways. First, we propose a new scheme for handling the technological complexity within a generic network planning approach. Second, we present computational studies for a realistic, large-scale indoor network installation. The basic idea of the new approach is to *combine* mathematical optimization and detailed network simulations. The complexity of the planning campaign of a large-scale multi-floor WLAN highly depends on the performance model selected. If the model used for optimization tries to account for the details of the CSMA-CA protocol as well as for interference effects, it can easily happen that the network planning problem turns out to be computationally infeasible. We propose to off-load complexity from the optimization problem and shift it to the simulation stage. For instance, protocol issues are exclusively tackled by the network simulation, while the optimization focuses on simplified metrics accounting for coverage and interference.

Our approach consists of three ingredients. First, we extend the two network performance criteria from [12] while keeping them simple enough to be susceptible to mathematical optimization. While *average net rate* depends only on the average link quality, *co-channel overlap* quantifies the interference potential of APs operating on the same channel. These aspects are partially contradicting. To maximize average net rate, we would densely cover the area with APs. To minimize co-channel overlap, we would seek to position APs transmitting on the same channel as far apart from each other as possible. Second, we apply multi-criteria optimization techniques to identify a set of Pareto optimal design candidates. Technically, we use an ϵ -constraint scalarization technique, in which different parameter choices produce distinct *Pareto optimal* network designs. Third, using a refined simulation model we make a realistic performance assessment of selected Pareto optimal designs. The simplified performance metrics of the optimization model prove sufficient to drive the designs towards the desired performance trade-off. Yet, there are quantitative and qualitative differences among the designs that are not traced using the simplified metrics.

Hence, these three ingredients together allow a balanced handling of the trade-off between performance modeling detail and duration of a planning campaign in the specific context of WLANs. This constitutes an effective approach for WLAN

planning, which is new to the best of our knowledge. Neither of [7]–[12] have the second or third ingredient.

The paper is organized as follows. Sec. II recalls the basic properties of IEEE 802.11 relevant to this work and contains our planning assumptions. Sec. III states two basic optimization models as well as their combination in the context of multi-criteria optimization. Sec. IV presents our simulation approach. Sec. V gives the optimization and simulation setups as well as the obtained results. Sec. VI contains conclusions and suggestions for future research.

II. WLAN PERFORMANCE AND PLANNING

The IEEE 802.11 standard [13] dominates among deployed WLANs. We first give a brief overview of their physical (PHY) and medium access control (MAC) layer schemes. We also sketch the key drivers of multi-cell WLAN performance and discuss how network planning affects them.

A. The IEEE 802.11 Standard

The standard defines different PHY and MAC schemes in a series of amendments. From these, version 802.11g is currently most widely deployed and at the focus of this article. We only consider the *infrastructure mode*, where stations communicate via *access points* (AP). For more details we refer to [13].

WLANs commonly organize medium access with the *distributed coordination function* (DCF). The DCF implements CSMA-CA protocol with binary exponential backoff, in which the AP behaves like a normal station. Basically, every station that has data to be transmitted senses the channel first and transmits when the channel is free for a certain time period. A successful data transmission is indicated by an acknowledgment. If the station does not receive such an acknowledgment, it concludes that a collision occurred and initiates the exponential backoff procedure. Optionally, a data transmission can be preceded by an RTS/CTS handshake, which reserves the channel during the data transmission by distributing a *network allocation vector* (NAV). Any station that senses the channel, checks this NAV in addition to comparing the currently received power on the channel with a *clear channel assessment* (CCA) threshold $\theta^{(c)}$.

Once a station went through the contention phase successfully, the data frame is transmitted using one out of twelve available PHY rates. This PHY data rate of the payload frame can attain up to 54 Mbit/s and is usually selected dynamically by the wireless network card, referred to as rate adaptation (not being part of the standard). IEEE 802.11g operates in the unlicensed band at 2.4 GHz. The standard divides the 2.4 GHz band into up to 14 different channels, while channel availability varies with regional regulations. Each channel has a 20 MHz frequency bandwidth and adjacent channels are 5 MHz apart. Hence, channels do partly overlap. We adopt the common assumption that there are exclusively three non-overlapping channels, namely 1, 6 and 13. In typical commercial deployments, each AP transmits on a single channel, which is assigned a priori.

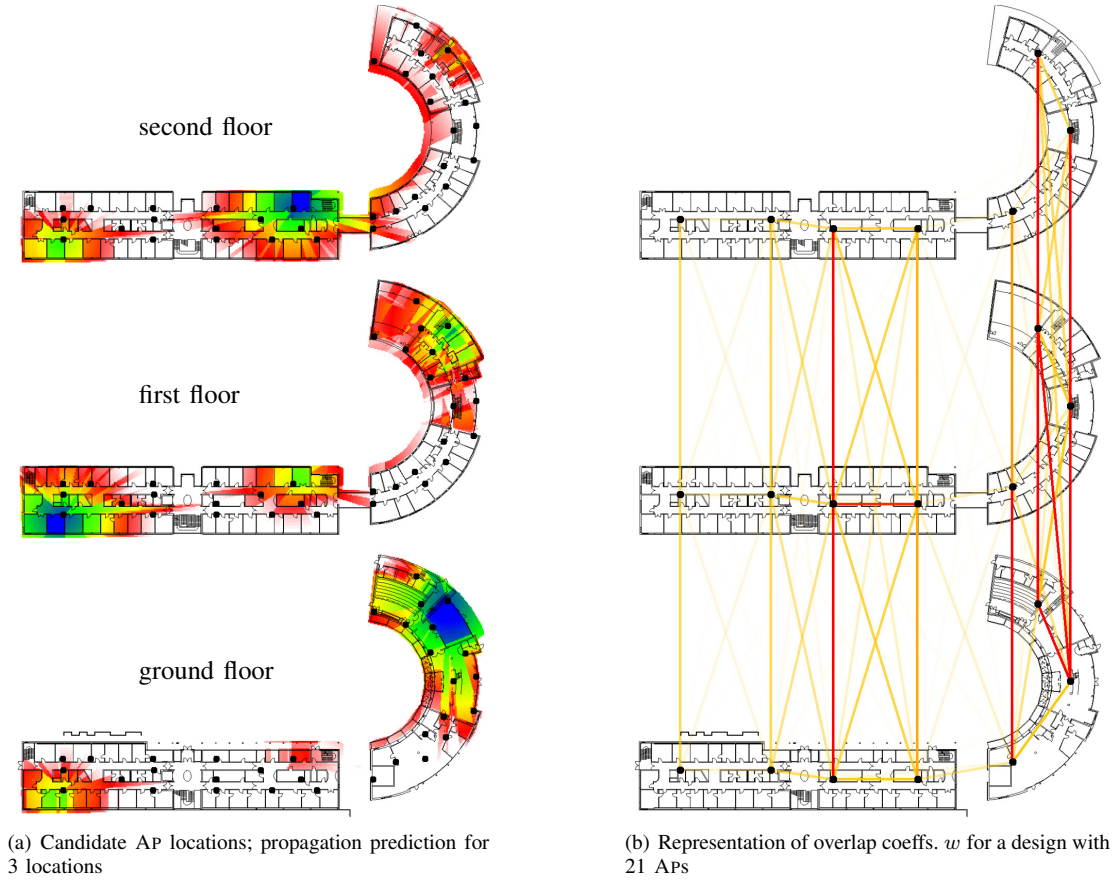


Fig. 1: Input data and simplified performance measures

B. WLAN Performance: Metrics, Factors and Optimization

The goal of wireless network planning is to optimize performance. WLAN performance can be judged by several metrics, e. g., throughput, goodput, delay, jitter, signal-to-noise ratio. The appropriate choice depends on the application. For a single cell, the physical data rate and the effects of contention under CSMA-CA are the two main factors that drive network performance. The stronger a station's signal is received, the higher is the potential physical data rate on this link. However, any signal that is received below the receiver sensitivity $\theta^{(s)}$ cannot be decoded. Accordingly, a *transmission range* can be defined around the position of a transmitting node.

When several stations access the network at the same time, the effects of contention kick in. Stations within *carrier sensing range* block each other when using the medium, where this range is determined by the CCA threshold (the larger the threshold, the smaller is the range). Also, WLAN features a *performance anomaly* due to rate adaptation [14]. High-rate stations experience essentially a much lower average net rate if one or several low-rate stations are also served by the AP.

Interference can have a strong influence on network performance. If the total cross-cell interference is below the CCA threshold, interference increases the packet error rate of the link. This can interact with the rate adaptation algorithm. If the cross-cell interference is above the CCA threshold, data transmissions in neighboring cells lead to pausing the countdown

of the backoff for all stations that are within the carrier sensing range. Furthermore, these stations might defer channel access or simply do not reply a RTS frame as the channel is sensed busy. These facts may degrade the performance significantly, but are hard to model mathematically.

WLAN performance is improved in several ways. The standard itself is reviewed and extended by amendments that alter the PHY and MAC protocol (providing more bandwidth in the PHY, introducing frame aggregation in the MAC, *etc.*). Next, an operating WLAN can adapt protocol parameters such as the CCA threshold, the applied rates per link, the transmit power, *etc.*, to the channel in order to improve network performance in an *online* fashion. Moreover, a WLAN deployment can be optimized prior to operation as part of network planning. This is an *offline* approach. In network planning the deployment of the infrastructure has to be decided, *i. e.*, in case of a WLAN this refers to the placement of the APs and their configuration. The configuration can consist of several parameters such as the channel assignment, the transmit power setting, the usage of specialized antennas. We refer to a particular decision regarding these variables as a *network design*. The objective in network planning is to determine a good network design.

C. Planning Assumptions for this Study

We choose as our primary measure of network performance the average *goodput* per station, *i. e.*, the data rate that the




Fig. 2: Physical data rate and station data rate in 802.11g

network provides on average per station above the link layer (in uplink and downlink direction). Instead of focusing on particular applications in order to model the traffic, we consider the *saturation mode* as load model, which is characterized by full buffer states in both up- and downlink at all times. We focus on WLAN deployments consisting of stations and APs that comply to the IEEE 802.11g amendment (with the exception that we only consider the eight OFDM-based PHY rates). Furthermore, we focus on infrastructure mode.

We take two planning decisions: locations and radio channels for APs. We do not optimize the APs' transmit power, but take a fixed value. We assume that the network planner has previously determined a set of potential AP locations according to the topology of the planning area and practical constraints. For example, some areas may be avoided for security reasons while other areas are preferable for ease of maintenance. An example of 114 candidate AP locations for a three-floor office building is illustrated in Fig. 1(a).

III. FINDING PARETO OPTIMAL NETWORK DESIGNS

A system model is an outcome of a detailed system analysis using all engineering expertise available. The analysis develops a clear understanding of what is to be achieved, what are the important parameters, and what are the important quality indicators of a solution.

While a detailed system modeling of a network and its behavior is necessary for performance evaluation, such detailed system models are typically too complex for network optimization. One of the challenges in devising a powerful optimization approach is to identify alternative models that are accessible to refined mathematical optimization techniques. See [15] for several examples of this approach.

In the first step of our planning approach, we optimize for two simplified performance measures for the downlink: *average net rate* and *co-channel overlap*. We introduce separate models for placing APs as to maximize average net rate (Sec. III-A) and for assigning channels as to minimize co-channel overlap (Sec. III-B). These models are first introduced

TABLE I: Variables and coefficients in optimization models

Name	Domain	Interpretation
Sets		
\mathcal{A}		Set of candidate APs
\mathcal{T}		Set of TPs
\mathcal{AT}	$\subset \mathcal{A} \times \mathcal{T}$	Set of feasible assignments from APs to TPs
\mathcal{C}		Set of available channels that do not interfere
Variables		
z_a	$\{0, 1\}^{\mathcal{A}}$	Is AP a selected
x_{aj}	$\{0, 1\}^{\mathcal{AT}}$	Is AP a assigned to TP j
y_{ab}	$\{0, 1\}^{\binom{\mathcal{A}}{2}}$	Do APs a and b operate on the same channel
f_a^c	$\{0, 1\}^{\mathcal{A} \times \mathcal{C}}$	Does AP a operate on channel c
Coefficients		
$p_{aj}^{(RX)}$	≥ 0	Received signal power from AP a at TP j
$\phi(p_{aj}^{(RX)})$	≥ 0	Net rate from AP a to TP j
N	> 0	Maximal number of APs to be selected
$\theta^{(C)}$	> 0	CCA threshold
$\theta^{(S)}$	> 0	receiver sensitivity
w_{ab}	≥ 0	Approximated overlap between APs a and b

in [12], where they are also discussed in detail. The notion of co-channel overlap used here, however, is different and better describes the potential up- as well as downlink interference.

This is an application of *multi-criteria optimization* [16], where a vector of objectives ought to be maximized. Since vectors cannot be compared like scalar numbers, the notion of optimality is different. A solution is said to be *optimal* (or *non-dominated*), if none of the objective function values can be improved without deteriorating any other objective function value. The optimization solutions with this property are called the *Pareto frontier* of the optimization problem. To solve the optimization problem, we ideally have to identify the Pareto frontier. A decision maker can then select the member of the Pareto frontier that seems most favorable. In Sec. III-C, we thus blend the two models into one parameterized model.

We use the notion of a *test point* (TP) as a location within the building, where coverage and interference is monitored. Test points are picked every two meters in areas where WLAN stations are likely to be used. TP j is said to be covered from AP a if it is within the transmission range of a (i.e., the received signal strength $p_{aj}^{(RX)}$ is at least $\theta^{(S)}$). A TP j is said to be within the carrier sensing range of an AP or another TP if the received signal strength $p_{aj}^{(RX)}$ at j of their transmissions is above the CCA threshold $\theta^{(C)}$. The CCA threshold defines the minimum amount of power at the receiver in order for the station to consider the channel as busy. Any station that senses a busy channel refrains from accessing the medium. Tab. I gives an overview of the coefficients and variables used in the models below.

A. Maximizing Average Net Rate

The (downlink) average net rate reflects network coverage and the impact of automatic rate selection. The goal in net rate maximization is to optimize the transmission rates that the TPs achieve by positioning APs appropriately. The average net rate of a given network design therefore depends on the signal propagation properties of the installed APs. Signal propagation

for three candidate APs is indicated in Fig. 1(a).

Basically, the “closer” a station is to its serving AP, the higher is the signal-to-noise ratio (SNR). Assuming a given noise level (and potentially a fixed, constant interference level), a monotonically increasing function $\phi(\cdot)$ maps the received power $p_{aj}^{(\text{RX})}$ at TP j from AP a to the effective net rate experienced at j , see the smooth curve in Fig. 2. In a deployment, the interference level will typically vary over time and with the location. Therefore, the described relation between SNR and net rate is only indicative. Moreover, the mapping does take framing overhead from the protocol into account but not the effects of contention.

The model contains two types of variables. The binary selection variables z_a are set to 1 if the AP a is selected and 0 otherwise. To assign TPs to APs, binary assignment variables x_{aj} are included. The model sets x_{aj} to 1 if TP j is assigned to AP a . Only such assignments x_{aj} are considered within the set \mathcal{AJ} , where AP a covers TP j .

$$\max \quad \frac{1}{|\mathcal{J}|} \sum_{(a,j) \in \mathcal{AJ}} \phi(p_{aj}^{(\text{RX})}) x_{aj} \quad (1a)$$

$$\text{s. t.} \quad x_{aj} \leq z_a \quad \forall (a,j) \in \mathcal{AJ} \quad (1b)$$

$$\sum_{a:(a,j) \in \mathcal{AJ}} x_{aj} \leq 1 \quad \forall j \in \mathcal{J} \quad (1c)$$

$$\sum_{a \in \mathcal{A}} z_a = N \quad (1d)$$

$$z \in \{0, 1\}^{\mathcal{A}}, \quad x \in \{0, 1\}^{\mathcal{AJ}}$$

The objective (1a) is to maximize the average net rate, which is expressed by means of a weighted sum over all potential assignments x_{aj} of AP a to TP j and the respective rates $\phi(p_{aj}^{(\text{RX})})$. Constraints (1b) and (1c) guarantee feasible assignments by ensuring that selected APs are assigned to TPs and that at most one AP is assigned to every TP, respectively. Constraint (1d) fixes the number of selected APs to N .

B. Minimizing Co-Channel Overlap

The notion of co-channel overlap addresses the contention between stations associated to distinct APs that use the same channel. The overlap is not considered harmful, if the two APs are assigned different channels. Contention with stations from other APs lowers the average net rate, so the total overlap of cells using the same channel should be minimized.

Contention is outside the scope of static models, which are typically employed for optimization. We therefore use interference as a surrogate measure for possible contention in the up- as well as the downlink. The following definitions estimate the areas that may be affected by co-channel interference. This is then used to model co-channel overlap. The optimization model at the end of the section asks for assigning channels to APs such that co-channel interference (in this terms) is minimized. The *service area* of an AP comprises all TPs where the AP provides sufficient signal power to establish a link, *i. e.*, all TPs that are within transmission range of the AP. We measure the size s_a of the service area of AP a by:

$$s_a := \left| \left\{ j \in \mathcal{J} \mid p_{aj}^{(\text{RX})} \geq \theta^{(\text{s})} \right\} \right|$$

Service from AP a to an associated TP j may be impaired by an AP b if j is within the carrier sensing range of b (and a

and b share the same channel). The number of all TPs in the service area of a and within the carrier sensing range of b is:

$$w_{ab}^c := \left| \left\{ j \in \mathcal{J} \mid p_{aj}^{(\text{RX})} \geq \theta^{(\text{s})} \text{ and } p_{bj}^{(\text{RX})} \geq \theta^{(\text{c})} \right\} \right|$$

This quantity will typically over-estimate conflicts between APs a and b , because not all of the TPs within the service area of a are necessarily serviced from a .

Contention may also arise between two stations under the following conditions: The stations are at TPs j and k and within carrier sensing range of each other. Moreover, they are associated with different APs a and b that use the same channel (*i. e.*, the stations compete for the same medium). Finally, the station at TPs j is outside the carrier sensing range of AP b . Then, AP a and a station at j block the channel for data reception at k , because the station hears the RTS/CTS frame, but its AP b does not. This can also lead to an unsuccessful attempt of b to transmit data to k . This situation is equivalent to the exposed terminal problem. We capture this by:

$$w_{ab}^u := \left| \left\{ (j, k) \in \mathcal{J} \times \mathcal{J} \mid p_{aj}^{(\text{RX})} \geq \theta^{(\text{s})} \text{ and } p_{bk}^{(\text{RX})} \geq \theta^{(\text{s})} \right. \right. \\ \left. \left. \text{and } p_{jk}^{(\text{RX})} \geq \theta^{(\text{c})} \text{ and } p_{bj}^{(\text{RX})} < \theta^{(\text{c})} \right\} \right|$$

The definition of overlap coefficients for pairs of APs to be used within the optimization is:

$$w_{ab} := w_{ab}^c s_b + w_{ba}^c s_a + w_{ab}^u + w_{ba}^u$$

Fig. 1(b) illustrates the overlap coefficients by means of colored lines. The brighter and darker the color of a line is, the larger is the value of the corresponding coefficient.

To minimize approximated overlap, we seek to find a design that minimizes the sum over all overlap coefficients w_{ab} between APs a and b using the same channel. Binary decision variables y_{ab} are set to 1 if a and b use the same channel and to 0 otherwise. The optimization problem is as follows:

$$\min \quad \sum_{ab \in \mathcal{A} \times \mathcal{A}} w_{ab} y_{ab} \quad (2a)$$

$$\text{s. t.} \quad \sum_{\{a,b\} \subset H} y_{ab} \geq 1 \quad \forall H \in \binom{\mathcal{A}}{4} \quad (2b)$$

$$y_{ab} + y_{bc} \leq 1 + y_{ac} \quad \forall (a, b, c) \in \mathcal{A}^3 \quad (2c) \\ y \in \{0, 1\}^{\binom{\mathcal{A}}{2}}$$

Three non-interfering channels are assumed to be available. Thus, inequality (2b) ensures that among any 4 APs at least two operate on the same channel. Constraint (2c) is a version of the triangle inequality. By forcing y_{ac} to 1, if y_{ab} and y_{bc} are set to 1, this ensures consistent channel assignments.

C. Integrated Multi-Criteria Optimization

In practice, multi-criteria optimization problems are commonly solved using *scalarization techniques*, which cast multi-criteria problems into parameterized single-objective problems. For each parameter setting, an optimal solution of the single-objective problem represents a member of the Pareto frontier. Different members of the Pareto frontier are found by varying the parameters.

We use the popular ϵ -constraint method for scalarization. The minimization of co-channel overlap is used as objective function, while average net rate maximization is transformed into a constraint. Suppose the maximum and minimum average net rate in the optimization space are given by r and R , respectively. We introduce a lower *rate limit* through the parameter $\tau \in [0, 1]$ and solve the optimization problem:

$$\min \sum_{(a,b) \in \binom{\mathcal{A}}{2}} w_{ab} y_{ab} \quad (3a)$$

$$\text{s. t. } \frac{1}{|\mathcal{J}|} \sum_{(a,j) \in \mathcal{AJ}} \phi(p_{aj}^{(\text{RX})}) x_{aj} \geq r + \tau(R - r) \quad (3b)$$

$$f_a^c + f_b^c \leq 1 + y_{ab} \quad \forall a, b \in \mathcal{A}, c \in \mathcal{C} \quad (3c)$$

$$\sum_{c \in \mathcal{C}} f_a^c = z_a \quad \forall a \in \mathcal{A} \quad (3d)$$

$$x_{aj} \leq z_a \quad \forall (a, j) \in \mathcal{AJ} \quad (1b)$$

$$\sum_{a: (a,j) \in \mathcal{AJ}} x_{aj} \leq 1 \quad \forall j \in \mathcal{J} \quad (1c)$$

$$\sum_{a \in \mathcal{A}} z_a = N \quad (1d)$$

$$x \in \{0, 1\}^{\mathcal{AJ}}, y \in \{0, 1\}^{\binom{\mathcal{A}}{2}}, z \in \{0, 1\}^{\mathcal{A}}, f \in \{0, 1\}^{\mathcal{A} \times \mathcal{C}}$$

An explicit assignment of colors is used. Binary decision variables f_a^c decide whether AP a operates on channel c , in which case f_a^c is set to 1.

Constraint (3b) enforces the average net rate to be at least the rate limit, which lies between r and R and is defined by τ . Constraint (3c) forces the overlap variable y_{ab} to be set to 1 if APs a and b work on the same frequency c . Constraint (3d) says that every selected AP has to be assigned one frequency. As before, constraints (1c) and (1d) ensure consistent channel assignments and (1b) calls for the selection of the desired number of APs. The value r is determined by solving the above model without constraint (3b), while R is determined by the model from Sec. III-A.

The model (3) describes an \mathcal{NP} -hard optimization problem. Assigning three colors to the nodes of a graph such that the weights of all monochromatic edges are minimized constitutes an \mathcal{NP} -hard optimization problem [17]. This problem can be reduced to model (3) by a proper choice of the parameters.

IV. SIMULATING PARETO OPTIMAL NETWORK DESIGNS

The multi-criteria approach described in the previous section takes average net rate and co-channel overlap into account. Due to the underlying CSMA-CA protocol complexity it is unclear which trade-off between the two gives the best performance for the specific planning situation. This question can be answered either by experimental measurements or by simulation. Due to cost, effort, and reproducibility, we evaluate this question through detailed simulation.

Evaluating WLAN performance by means of simulation is a common practice throughout the community. WLAN simulation models are available for commercial simulation tools as well as for open source tools. However, the degree of similarity to the standard and level of detail vary significantly. Hence, a lot of work on performance evaluation of WLANs focuses on adding more realism to simulation models, for example regarding modeling the CCA procedure [18], the frequency-selectivity of subcarrier gains [19], [20] or accurate path loss coefficients [21]).

We judge a network design by its *expected* performance, c.f. Sec. II-B. The expected performance is assessed by Monte Carlo simulations. We draw load instances (snapshots) according to a random distribution and simulate network operation in detail until the observed metrics converge statistically. A load instance is characterized by the placement of the stations. As we assume saturation mode, this effectively results in a load instance, which we also call in the following a snapshot. The random distribution of the station positions is sampled multiple times, i.e., we consider various different snapshots and aggregate the results. In the following, we discuss the simulation approach and model.

A. Evaluation Procedure

a) Snapshot Generation: We generate a set of I different snapshots, characterized by J stations distributed with uniform probability over the considered area (more specifically: over the considered TPs). Given these positions, we associate stations to APs (from the respective design) according to the received signal strength, which is the standard AP selection procedure of IEEE 802.11 NICs. Received signal strength calculations are based on the same propagation model as used for the design generation (see Sec. V-A). Obviously, the same snapshot (i.e., the same distribution of the J stations over the considered area) may result in different station associations for different network designs. The traffic generated by all stations and APs is set to saturation mode, parameterized by a fixed packet size ς .

b) Performance Metrics: Given a snapshot, we simulate packet transmissions in the network for the different designs. After each simulation run we obtain the average MAC throughput, goodput and packet loss rate in up- and downlink as well as the PHY rate setting per station. As there is no single event that can quantify the exact impact of external interference on the performance of each cell, we conduct a second simulation run where all cells are artificially *isolated* from each other (switching-off cross-cell interference). During this run we maintain the snapshot and propagation characteristics within each cell. For C different network designs and I different snapshots, we conduct $C \cdot I$ simulation runs (per evaluation setting). We average the collected data per design (also calculating 95% confidence intervals) and compare the performance of different designs.

B. Simulation Model

Packet transmission is simulated according to standard IEEE 802.11 procedures regarding MAC and the upper part of PHY. For our simulations we use the 802.11 model of OPNET Modeler version 14.0.A-PL2. This is a state-of-the-art network simulator, offering a high level of detail in the modeling of the CSMA-CA protocol. Furthermore, we add frequency-selective fading to the existing simulation model as this has a significant impact on WLAN performance [19]. Channel gain coefficients consist of a path loss component determined individually for each link according to the model referred to in Sec. V-A. On top of this, a random frequency-selective fading component is added. For each subcarrier and each packet transmission we

TABLE II: Scenario characteristics

Category	Characteristic / Parameter	Value
Building	Floors	3
	Area / Floor	2 700 m ²
	Candidate AP locations	114
	APs	21
AP	Type	IEEE 802.11g
	Antenna gain	2 dBi
	Transmit power	30 mW
	Height above floor	2.2 m
Channel	Frequency band	2.4 GHz
	Channel set	13 channels (ETSI)
TP	Grid resolution	2 m
	Tps	1965
	Height above floor	1 m
Powers / Thresholds	Energy detection	−115 dBm
	RX noise power	−100 dBm
	Receiver sensitivity $\theta^{(s)}$	−95 dBm
	CCA threshold $\theta^{(c)}$	−95 dBm

draw a random coefficient from an exponential distribution. Based on this fading profile, the SINR per subcarrier and ultimately the bit error rate per subcarrier are generated. Using an analytical approach [20], we compute the equivalent packet error rate and drop the packet with this probability.

Rate adaptation is another issue with a strong impact on performance. We do not assume a certain algorithm to be in place. Based on the received signal strength (excluding fading), we set the rate for each station at the beginning of the simulation to a common fixed value for up- and downlink. The exact thresholds are depicted in Fig. 2. They stem from extensive simulation runs to determine the optimal settings with respect to goodput above the MAC (depending also on the activation of RTS/CTS handshake and the packet length). Any control packet transmitted during the simulation is conveyed with BPSK rate 1/2. No Beacon packets are transmitted. Mobility as well as station appearance/disappearance/reassociation are not taken into account.

V. COMPUTATIONAL STUDY

We test our approach on a large scenario with realistic data.

A. Planning Scenario

The planning scenario is based on the three-floor ZIB office building. A network planner has previously determined a set of potential AP locations according to the topology of the planning area and practical constraints. Out of a set of 114 candidate AP locations depicted in Fig. 1(a), 21 APs are to be chosen. Each AP is of type IEEE 802.11g and equipped with an omni-directional antenna. The 1965 TPs for throughput optimization are placed with a resolution of 2 m. Tab. II provides basic information on the planning scenario.

A *Reference* design for the scenario is known. This is the result of careful, manual configuration, which was established in several weeks of work. The *Reference* design is the one implemented within the building.

Radio propagation is predicted using the COST 231 multi-wall model [22, Ch. 4] with the empirical reduction of multi-wall effects as suggested in [23]. The model assumes that reflections can be neglected and only the distance and the obstacles on the direct path between the two points have to be taken into account. The obstacles between two points are determined in a 3D model of the building. The parameters of the model are set to produce predictions that match the measured path losses reasonably well.

B. First Step: Trading Off Net Rate vs. Co-Channel Overlap

We solve the integrated planning model from Sec. III-C using integer programming techniques (as implemented in commercial solver CPLEX, V11.0.0). Even though the optimization problem is \mathcal{NP} -hard in general, modern mixed integer programming solvers manage to solve the problem by means of Branch&Bound techniques with a bounding based on underlying LP relaxations. The optimal solution of our instances is usually found in less than 10 hours running time.

Different settings for τ are used to determine solutions along the Pareto frontier trading off average net rate vs. co-channel overlap. We employ a “divide & conquer” strategy. The building includes multiple floors and two separate parts, which are connected by a gangway, see Fig. 1. The problem is first solved separately for the rectangular part of the building on the left-hand side and the circular part on the right-hand side. The split is chosen in this way for the small interference coupling between those two parts of the building. Minor overlap is generated between those parts as well, which cannot be taken into account during the separate optimizations. After the AP locations of the two solutions are joined, the channel assignment is optimized according to Sec. III-B for the entire building in order to minimize the impact of pasting. Nevertheless, some disturbances of monotonicity in co-channel overlap remain for solutions with small values of τ .

As the requirements for average net rate increase, the network designs evolve from clusters of three APs using distinct channels with large spaces between them ($\tau = 0$) to an interleaved structure of APs ($\tau = 1$). The trade-off between co-channel overlap and average net rate is plotted in Fig. 3. The designs obtained by the integrated optimization are labeled with the corresponding values for τ . In [12] has been observed that a minor deterioration in co-channel overlap can be traded in for a large improvement in average net rate. For our eleven Pareto-optimal designs, when focusing on co-channel overlap only the area without coverage is largest (close to 6%). This loss then declines, and for several intermediate values of τ the loss is below 0.5%. Surprisingly, focusing on average net rate alone again yields a higher loss of about 3%.

C. Second Step: Simulation

After identifying a set of network designs with different relations between average net rate and co-channel overlap, the question remains what this means for actual station goodput and which design is best. This question shall be answered by detailed simulation. We would like to point out, that the performance of the selected network designs is closely

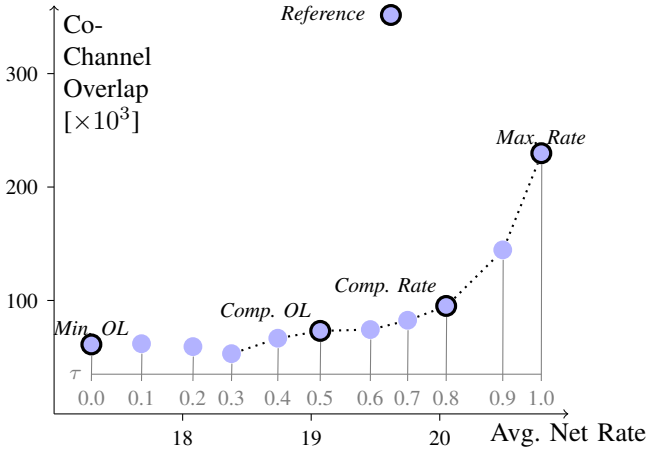


Fig. 3: The circled designs are selected for simulation. The dotted line (starting at 0.3) indicates the Pareto frontier w.r.t. the simplified performance measures. The average net rate in Mbps from Sec. III-A is shown on the x-axis, the spatial measure for co-channel overlap from Sec. III-B on the y-axis.

related to the building under consideration and its propagation characteristics, as well as to the considered network load. We are not proposing a certain network design as a sort of *universal* solution, since the performance of the network designs may be different in other scenarios and under different load conditions.

In our simulations we pick the following five designs highlighted in Fig. 3 for a close analysis:

- Simulation results for the *Reference* solution provide the baseline performance for the network design implemented. Recall that this design is the result of a careful, manual planning process.
- Two *compromise* solutions, *Comp. OL* tending towards co-channel overlap minimization ($\tau = 0.5$) and *Comp. Rate* towards average net rate maximization ($\tau = 0.8$), are our candidates for the best network designs.
- Two *extreme* solutions, *Min. OL* ($\tau = 0$) and *Max. Rate* ($\tau = 1$), are used to study the relationship of co-channel overlap and average net rate.

Each design is analyzed in four settings obtained by combining the following options:

- *Without* and *with* cross-cell interference, denoted by *isolated* and *interfered*, respectively.
- In the presence of pure downlink traffic or with up- as well as downlink traffic.

D. Results

We first look at the performance with pure downlink traffic. Fig. 4 presents the corresponding average goodput results per station. In the interfered case, we observe significantly different results (varying by about 30%) among the designs obtained by integrated optimization. The *Min. OL* design performs worst, while the *Max. Rate* design provides the highest goodput. Close to the performance of the *Max. Rate* design are the performances of the *Comp. OL* and *Comp. Rate*

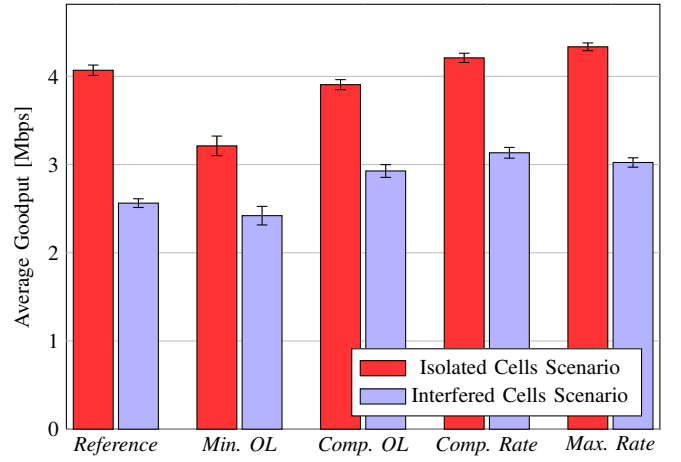


Fig. 4: Interfered and isolated average goodput for five network designs with *pure downlink* transmissions.

designs. The analysis without cross-cell interference allows to fathom how well each individual AP is servicing its stations. If cross-cell interference is ignored (for the *isolated* scenario), the *Max. Rate* design still provides best performance according to the Pareto frontier in Fig. 3. The designs that provide good means to fight interference (*Min. OL*, *Comp. OL*) are the ones that profit less from this artificial *decoupling* of the cells. In general, a significant amount of performance is lost due to interference, e.g., in the case of *Comp. Rate* this constitutes 30% of the goodput.

Furthermore, we observe that the *Reference* design provides good coverage. This is highly desirable, as this design is implemented in practice. Although extensive manual planning has produced good coverage results here, the high effort is not a desirable option in practice. It is not only costly in terms of time (several weeks of work) but it is also subject to performance uncertainty. This is strengthened by the fact that the *Reference* design is clearly outperformed by automatic planning in the presented results. In a second round, we consider up- and downlink transmissions (Fig. 5). Recall that IEEE 802.11 uses a stochastic form of time division duplex (TDD), therefore up- and downlink transmissions compete for the same medium. Per cell, only one AP transmits in the downlink, while typically a much higher number of stations transmit in the uplink. Hence, the downlink goodput observed in the pure downlink setting drops considerably in the presence of balanced up- and downlink demands. This explains the smaller values in downlink goodput and the much higher values in the uplink case (e.g., in the isolated case the difference in performance is about 400%).

Focusing on the downlink in the interfered case (Fig. 5(a)), the *Max. Rate* design provides the best performance, followed by the *Comp. Rate* design. All other designs provide almost the same downlink performance, which is about 30% lower than for *Max. Rate*. Again, all designs suffer significantly from interference as we observe for all five designs large performance gains in the isolated scenario.

Finally, notice the different performance results for the uplink. Let us first focus on the interfering results. Here,

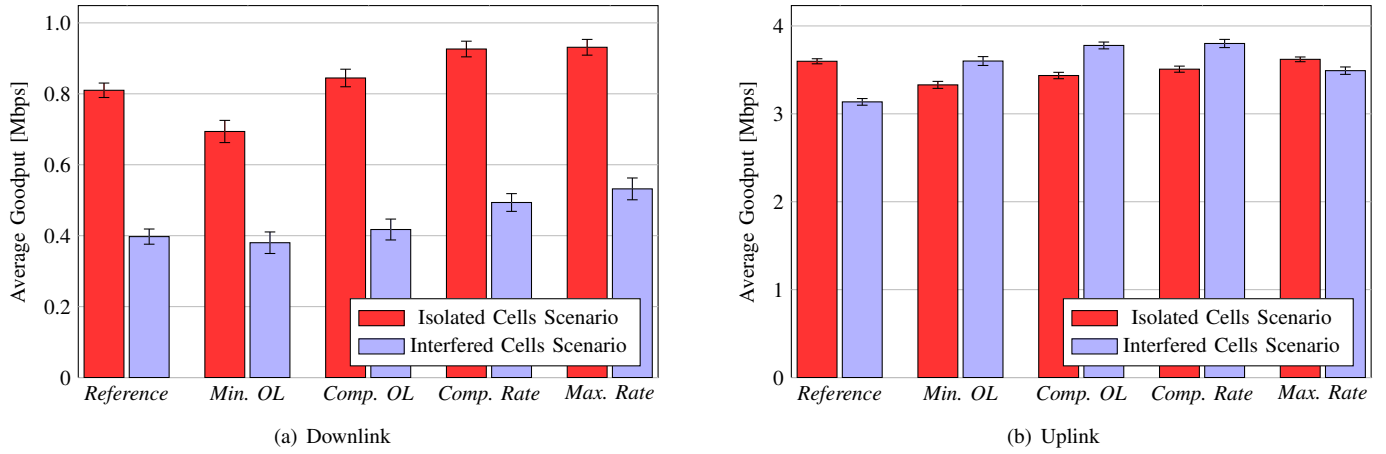


Fig. 5: Interfered and isolated average goodput for five network designs with *up-* and *downlink* transmissions.

Max. Rate provides the worst performance among the designs from integrated optimization, *Comp. Rate* performs best and the *Reference* design is the worst among all designs. Furthermore, if no cross-cell interference is taken into account as reflected in the isolated case, uplink performance *decreases* significantly for the designs *Min. OL*, *Comp. OL* and *Comp. Rate*. This is due to the strong increase in downlink goodput when switching-off interference. As the AP is significantly more successful in transmitting packets per channel access (under RTS/CTS handshake), the time spent as part of the exponential backoff procedure is reduced. Consequently, the AP contends more often for the channel (the corresponding results are not shown here due to the lack of space). This leads to less remaining air time that the uplink transmissions are contending for. As interference plays a minor role in the data reception at the access point, the isolation is not improving the transmission situation physically for the uplink connections. Again, notice that all these interactions between uplink and downlink performance as well as the specialties of the CSMA-CA protocol cannot be included into the optimization problem and their impact on the performance cannot be foreknown. Simulations serve that purpose and fill the existing gap in the planning campaign.

The previous results present average values for the goodput. However, they do not reflect where the performance comes from, *i.e.*, how the goodput is distributed among stations (and APs). To investigate this issue, another set of results is presented. Figure 6 shows the cumulative distribution function of the goodput per station for the downlink (left) and the uplink (right), respectively. In the downlink, the *Min. OL* design clearly performs worst. The 50-th percentile ($CDF = 0.5$) of this design indicates that half of the stations have a goodput below 0.1 Mbps. We also observe that there is a significant amount of stations with an extremely low goodput (10^4 bps or lower). The other four designs exhibit much better performance at this point. This feature is closely related to the fact that the *Min. OL* design focuses on the reduction of interference, thus deteriorating the coverage and the achievable average net rate of numerous stations. *Max. Rate* exhibits the best performance, since it is able to keep the proportion

of *low*-goodput stations below the one achieved by most of the other designs, while also achieving a high proportion of *high*-goodput stations. The good coverage provided by the *Reference* design is corroborated by the fact that this design has the lowest proportion of *low*-goodput stations. In the uplink, the goodput achieved is significantly higher than in the downlink. It is important to note, that the performance gap between the *Min. OL* design and the best performing designs (*Comp. OL* and *Comp. Rate*) is not as significant as in the downlink case. It is in the cell edges where cross-cell interference causes a higher damage. Therefore, the stations and not the APs are the nodes that suffer most from interference. Correspondingly, uplink transmissions are hindered due to the CCA mechanism much more often than downlink ones. Hence, in the uplink case, the *Min. OL* design exhibits a considerably better behavior than in the downlink. However, this design is still far from yielding a good performance due to its lower average net rate. This is in line with the uplink performance of the *Max. Rate* and *Reference* designs, which perform significantly worse than the other designs over a vast CDF range. Both designs basically try to maximize the coverage and the average net rate regardless of the amount of potential interference. Such an approach is, in the uplink case (and at least for the scenario considered), definitely not the best choice.

VI. CONCLUSION

We propose a new method for dealing with the complex problem of large-scale multi-floor WLAN planning. Using multi-criteria optimization methods, we generate several network designs. These designs represent different trade-offs between co-channel overlap reduction and average net rate maximization, which are two simplified performance metrics. The designs are taken as candidates for the final planning solution. To pick the best design, we conduct a detailed analysis by simulation. This is tractable because only few candidate designs are left.

Our two-stage approach is superior to traditional planning techniques. Both mathematical optimization and detailed simulation can play to their strengths. Mathematical optimization

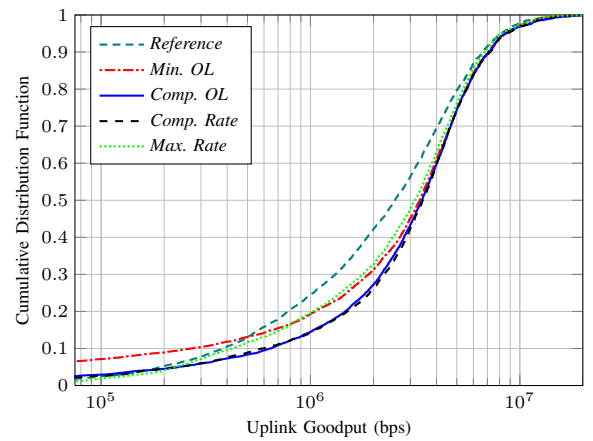
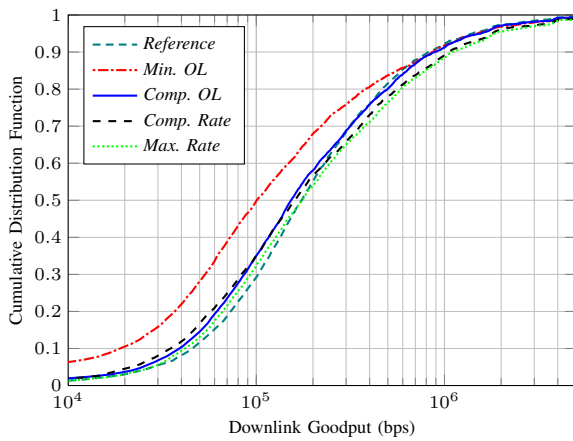


Fig. 6: CDFs of the average goodput per station in the downlink (left) and uplink (right) for five network designs

is able to find globally good designs according to simplified metrics. We employ it to fathom the general search direction. Because no single simple optimization objective is able to faithfully reflect WLAN performance, we use two competing measures. Detailed simulation is able to accurately compare the performance of designs. This takes time, and we thus use this tool only on the small set of sensible designs that mathematical optimization has identified. The combination of the two allows us to perform network planning, especially for large-scale multi-floor deployments, where either method has significant drawbacks. Our numerical results strengthen these claims. The chosen Pareto-optimal designs differ considerably in their performance which can also be observed in other planning scenarios.

For future work, we plan to investigate the relationship between pre-deployment network planning and resource management during network operations. Some authors argue that an intelligent radio resource management, for example, performed by dynamic station associations depending on AP loads, renders network planning virtually pointless. This assertion will be fathomed by comparing the detailed performance of well-planned and trivial network design with and without such dynamic resource management schemes.

REFERENCES

- [1] S. Hurley, "Planning effective cellular mobile radio networks," *IEEE Trans. Veh. Technol.*, vol. 12, no. 5, pp. 243–253, 2002.
- [2] R. Mathar and M. Schmeink, "Optimal base station positioning and channel assignment for 3G mobile networks by integer programming," *Ann. Oper. Res.*, no. 107, pp. 225–236, 2001.
- [3] M. Kamenetsky and M. Unbehauen, "Coverage planning for outdoor wireless LAN," in *Proc. IZS2002*, Zurich, Switzerland, 2002.
- [4] E. Amaldi, A. Capone, M. Cesana, L. Fratta, and F. Malucelli, "Algorithms for WLAN coverage planning," *Springer LNCS*, vol. 3427, pp. 52–65, Feb. 2005.
- [5] K. K. Leung and B.-J. Kim, "Frequency assignment for IEEE 802.11 wireless networks," in *Proc. VTC 2003-Fall*, Orlando, FL, 2003.
- [6] J. Riihijärvi, M. Petrova, and P. Mähönen, "Frequency allocation for WLANs using graph colouring techniques," in *Proc. WONS'05*, St. Moritz, Switzerland, 2005.
- [7] P. Wertz, M. Sauter, G. Wölfe, R. Hoppe, and F. Landstorfer, "Automatic optimization algorithms for the planning of wireless local area networks," in *Proc. VTC2004-Fall*, Los Angeles, CA, Oct. 2004.
- [8] R. C. Rodrigues, G. R. Mateus, and A. A. F. Loureiro, "On the design and capacity planning of a wireless local area network," in *Proc. NOMS2000*, Honolulu, Hawaii, 2000.
- [9] X. Ling and K. L. Yeung, "Joint access point placement and channel assignment for 802.11 wireless LANs," in *Proc. WCNC 2005*, New Orleans, LA, 2005.
- [10] T. Vanhatupa, M. Hannikainen, and T. Hamalainen, "Genetic algorithm to optimize node placement and configuration for wlan planning," in *Wireless Communication Systems, 2007. ISWCS 2007. 4th International Symposium on*, Oct. 2007, pp. 612–616.
- [11] A. Gondran, O. Baala, A. Caminada, and H. Mabel, "Joint optimization of access point placement and frequency assignment in wlan," in *Internet, 2007. ICI 2007. 3rd IEEE/IFIP International Conference in Central Asia on*, Sept. 2007, pp. 1–5.
- [12] A. Eisenblätter, H.-F. Geerdes, and I. Siomina, "Integrated access point placement and channel assignment for wireless LANs in an indoor office environment," in *Proc. WoWMoM'07*, Helsinki, Finland: IEEE, 2007.
- [13] IEEE Standards Association, "IEEE Std 802.11-2007," <http://standards.ieee.org/getieee802/802.11.html>, 2007.
- [14] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. INFOCOM 2003*, San Francisco, CA, 2003.
- [15] A. Eisenblätter and H.-F. Geerdes, "Wireless network design: Solution-oriented modeling and mathematical optimization," *IEEE Trans. Wirel. Comm.: 3G/4G/WLAN/WMAN Planning and Optimization*, pp. 8–14, Dec. 2006.
- [16] M. Ehrgott, *Multicriteria Optimization*, 2nd ed. Springer, 2005.
- [17] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi, *Complexity and Approximation: Combinatorial optimization problems and their approximability properties*. Springer, 1999.
- [18] H. Ma, E. Alotaibi, and S. Roy, "Analysis and Simulation Model of Physical Carrier Sensing in IEEE 802.11 Mesh Networks," in *Proc. OPNETWORK 2006*, August 2006.
- [19] O. Awoniyi and F. Tobagi, "Packet Error Rate in OFDM-based Wireless LANs Operating in Frequency Selective Channels," in *Proc. IEEE INFOCOM*, April 2006.
- [20] J. Gross, M. Emmelmann, O. Puñal, and A. Wolisz, "Dynamic single-user OFDM adaptation for IEEE 802.11 systems," in *Proc. ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWIM 2007)*, Chania, Crete Island, Oct. 2007, pp. 124–132.
- [21] A. Schmitz and M. Wenig, "The effect of the radio wave propagation model in mobile ad hoc networks," in *Proc. ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWIM 2006)*, Oct. 2006, pp. 61–67.
- [22] E. Damosso, "Digital mobile radio towards future generation systems—COST 231 final report," COST 231, Tech. Rep., 1999, available online at <http://www.lx.it.pt/cost231/>.
- [23] AWE Communications GmbH, "Empirical propagation models—fast empirical models for indoor scenarios," <http://www.awe-communications.com/Propagation/Indoor/Empirical>, retrieved on 19.05.2008.