TKN Telecommunication Networks Group

Technical University Berlin

Telecommunication Networks Group

Applying ad-hoc relaying to improve capacity, energy efficiency, and immission in infrastructure-based WLANs

M. Kubisch, S. Mengesha, D. Hollos, H. Karl, A. Wolisz

kubisch,mengesha,hollos,karl,wolisz@ee.tu-berlin.de

Berlin, July 2002

TKN Technical Report TKN-02-012

TKN Technical Reports Series

Editor: Prof. Dr.-Ing. Adam Wolisz

Abstract

In classical infrastructure-based wireless systems such as access point-equipped wireless LANs, all mobile terminals communicate directly with the access point. In order to improve the capacity or energy efficiency of such systems, the use of mobile terminals as intermediate relays has been proposed. The rationale is that intermediate relays reduce the communication distance and hence the emitted power. Therefore, relaying could also reduce electromagnetic immission. To assess these potential benefits, we study the effectiveness of various relaying algorithms in a uniform, HiperLAN/2-based system model that has been amended by relaying functionality. These algorithms jointly select intermediate relay terminals and assign transmission power as well as modulation to mobile terminals.

The energy efficiency of a point-to-point communication is indeed improved by relaying, however, this effect only marginally transfers to scenarios taking into account several terminals. Nevertheless, it is still possible to extend the lifetime of a network by taking into account available battery capacities.

For a discussion of capacity improvements, two different modes of conceiving system fairness are identified. For both system modes, we present relaying algorithms. Moreover, adding an additional frequency to a cell is beneficial: Using two frequencies can almost double the cell capacity, and for one fairness mode, even relaying with one frequency can improve capacity by up to 30 %.

In addition, all these algorithms reduce the immitted power averaged over the area of a cell. All algorithms show improvements, an additional algorithm specialized to reduce immission power can cut the average power almost in half.

Contents

Keywords: Relaying, capacity, energy efficiency, immission, WLAN, HiperLAN/2

Chapter 1

Introduction

The use of multi-hop wireless communication systems has been advocated where a sender communicates with its destination with the help of intermediate relay stations. Relaying is a necessary ingredient in realizing ad-hoc, spontaneously formed networks with no infrastructure support where not all the mobile terminals are in direct radio contact with each other. But relaying can also be applied in infrastructure-based, cellular communication systems in order to improve different aspects of their operation. Figure 1.1 shows such a scenario consisting of a far terminal A relaying via B to an access point.

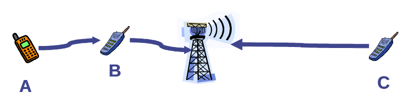

Figure 1.1: Example scenario for choosing relaying terminals

The expectation why relaying could actually be beneficial in an infrastructure-based environment is the reduction of distances over which communication takes place: No longer do all terminals have to communicate with a potentially far-away access point, but rather could communicate with an intermediate relay that is perhaps halfway towards the access point. Relaying allows the mobile terminal to reduce the power it uses for transmission. This power reduction can be leveraged to improve three key properties of a wireless communication system: energy efficiency, capacity, and electromagnetic immission.

Energy efficiency can be improved by the non-linear relationship of distance and transmission power: Communicating twice over short distances requires a smaller total radiated power than just once over a long distance. This could lead to a more efficient use of energy stored in a mobile terminal's batteries, but also requires to account for the additional overheads in receiving transmission, powering amplifiers, etc.

Capacity can be improved because smaller radiated power also means smaller interference. Shorter distances also allow the use of faster modulations. In addition, relaying could also open the way to add additional resources, i.e., frequencies, to a cell and thus increase cell capacity.

Electromagnetic immission can be improved because of the immediate benefits of using smaller transmission power levels, resulting in smaller average or peak immission values, depending on the distance from the access point.

Yet another aspect of relaying is the extension of the coverage of an access point by means of relaying. This paper focuses on relaying *within* a single cell, complementing other research focusing on coverage extension (e.g., [13]). While this restriction limits the problem scope, it allows to pursue simpler solutions for the relaying entities.

Such a relaying entity could be part of a fixed infrastructure (and connected to a fixed power supply), or the mobile terminals themselves could act as relaying entities whenever this is necessary and beneficial for the entire system. In the former case, much of the design decisions are quite simple and it is straightforward to see improvements with regard to these three aspects. The later case is much more challenging as the system is much more dynamic and as a relaying terminal, unlike a fixed relaying entity, also has its own traffic requirements and user needs that it has to support. We concentrate in this paper on the second case of purely using mobile terminals as relaying entities.

The contribution of this paper is to show how ad-hoc relaying via mobile entities can be integrated into a cellular system and to evaluate these three aspects in a uniform, mobile system environment. For each aspect we describe algorithms that determine which terminal is relaying data for which relayed terminal (essentially, solving a routing problem) and how these communication relationships are organized in time (in essence, a scheduling and modulation selection problem). In addition, assigning appropriate transmission power levels to mobile terminals is a crucial issue. As it will turn out, jointly optimizing transmission power and modulation schemes used for each communication is the key technique to leverage the potential benefits of relaying.

The remainder of this paper is organized as follows. Section 2 describes the model assumptions used here. Sections 3, 4, and 5 describe our algorithms for optimizing energy efficiency, capacity, and electromagnetic immissions, along with simulation-based evaluations of their performance.Finally, Section 6 concludes the paper.

Chapter 2

System model

While extending infrastructure-based systems by relaying can be researched in the context of classical cellular systems like GSM or UMTS, this paper concentrates on access point-based wireless LANs as basic technology. Currently, there are two viable candidates for WLAN systems: IEEE 802.11 a/b as well as HiperLAN/2. While IEEE 802.11 is more popular at the moment, HiperLAN/2 will be able to provide comparable bandwidth and a superior QoS support. In addition, its centralized system structure allows to easier experiment with different routing and scheduling algorithms: In HiperLAN/2, the access point is responsible for computing a communication schedule for a MAC frame 2 ms long. In a frame, each mobile terminal is assigned time slots in which it is allowed to send or receive data to or from the access point or other mobile terminals; moreover, this schedule also stipulates the transmission power a mobile uses and which one out of seven different, standardized modulation types is used within a time slot.[1]

We extended this frame structure by a relaying protocol that allows traffic to travel from a mobile terminal to the access point via an intermediate relaying terminal (and vice versa) [10]. This extension is very lightweight but, unlike other extensions [13], does not result in an extended coverage of the access point. Using this protocol, the access point can use different algorithms to compute transmission schedules for a MAC frame, depending on which performance metric is to be optimized. The algorithms work on the basis of the current channel gains between access point and mobile terminals (which are available with a certain precision using HiperLAN/2's radio map). In the following, we identify the channel gains with the relative positions of the terminals. As we do not consider any obstacles in our evaluation, this is an acceptable simplification.

The next three sections will present algorithms that are optimizing energy efficiency, capacity, and electromagnetic immission. These sections also present performance results, based on a uniform simulation model [9]: A single cell is considered, mobile terminals are randomly placed on a square area of 70 m x 70 m, the access point is in the middle. The path loss coefficient α for most of the simulations is set to 3.2; all confidence levels are 95 %.

[1]Nevertheless, we expect our results in principle to carry over to IEEE 802.11-based systems as well since a similar physical layer is used, albeit a practical implementation is going to be somewhat quite from an HiperLAN/2 system. The main difference the lack of a central control, necessitating a distributed version of our algorithms.

Chapter 3

Optimizing energy efficiency

In order to use relaying to optimize energy efficiency, first an appropriate model of energy consumption must be developed. In recent publications [3, 6] energy models for relaying with adaptive transmission power have been developed, but either energy necessary to receive data has been neglected or power consumption not relevant to energy efficiency (e.g., power consumption in the AP) has been considered.

The first step for such a model is to define the power consumed in a single terminal. The following model is based on measurements [2, 4] which indicate a transmission power consumption behavior that can be approximated as follows:

$$P_{\mathrm{tx}} = a \cdot P_{\mathrm{radiated}} + P_{\mathrm{txFix}}$$

where P_{tx} is the total power consumed while transmitting data, a is a proportionality factor representing the amplifier's power consumption, P_{radiated} is the network card's actual power output and P_{txFix} is the power needed for amplifier-unrelated parts, e.g., baseband processing. Both power required to receive data P_{rx} and idle power P_{idle} are assumed to be constant.

Choosing values for these constants is in principle straightforward, except for P_{idle}: "idle" could refer to a network card neither sending nor receiving, or it could mean a network card that is actually powered down to a sleep mode [14]. On the one hand, idle power in the first sense is not much different from P_{rx}; on the other hand, the tightly scheduled HiperLAN/2 structure easily allows to assume sleep modes for inactive wireless terminals. Hence, in the following evaluation, the P_{idle} value is chosen considerably smaller than P_{rx}. While this neglects energy necessary to switch between idle and sleep modes, this can be assumed to be a minor source of energy dissipation.

Sending a fixed amount of data can be done by using different modulations, requiring different amounts of time. In order to determine the modulation with the highest energy efficiency to be used between two terminals, a target packet error rate (PER) has to be set. Based on this target PER, the optimal radiated power can be computed as a function of the distance [12], and the average energy per *correct* bit can be computed. Hence, a mapping from target PER and distance to the optimal modulation and radiated power can be derived.

Using this mapping, the total energy required to transmit a fixed amount of data either by direct communication of a far source terminal $\mathrm{MT}_{\mathrm{FAR}}$ with the AP or by relaying via a relay candidate $\mathrm{MT}_{\mathrm{REL}}$ can be computed.[1] Here, the frame structure of HiperLAN/2 and the resulting idle times

[1] The target PER for the relay case is smaller than for the direct case in order to result in the same end-to-end PER.

have to be taken into account. The equations for these cases are:

$$\begin{aligned} E_{\text{direct}} =& P_{\text{tx,FAR}\rightarrow\text{AP},i} \cdot T_i + P_{\text{idle}} \cdot (T_{\text{frame}} - T_i) + \\ & P_{\text{tx,REL}\rightarrow\text{AP},k} \cdot T_k + P_{\text{idle}} \cdot (T_{\text{frame}} - T_k) \end{aligned} \quad (3.1)$$

$$\begin{aligned} E_{\text{relay}} =& (P_{\text{tx,FAR}\rightarrow\text{REL},k} + P_{\text{rx}}) \cdot T_k + P_{\text{idle}} \cdot (T_{\text{frame}} - T_k) + \\ & 2 \cdot P_{\text{tx,REL}\rightarrow\text{AP},k} \cdot T_k + P_{\text{idle}} \cdot (T_{\text{frame}} - (2 \cdot T_k)) \end{aligned} \quad (3.2)$$

where E_{direct} and E_{relay} are the total energy for either case, $P_{\text{tx},X\rightarrow Y,i}$ is the power necessary for transmitting at modulation i and T is the time for sending or receiving the number of used information units at modulation i. Equation (3.2) describes the case of a relay terminal which also has data of its own to transmit; it has to be modified for an idle relay or a relay terminal that already is relaying data for some other terminal.

As an example, Figure 3.1[2] shows the ratio between total energy consumed in the direct and relaying case when the relay terminal is placed halfway between far terminal and AP. Relaying can considerably improve energy efficiency, but only for rather large distances of relayed terminal and access point (already approaching coverage limits) and optimal relay position; a combination that does only rarely happen in practice. As expected, reducing the fixed power offset P_{txFix} offers more possibilities for relaying and improves the relaying gain.

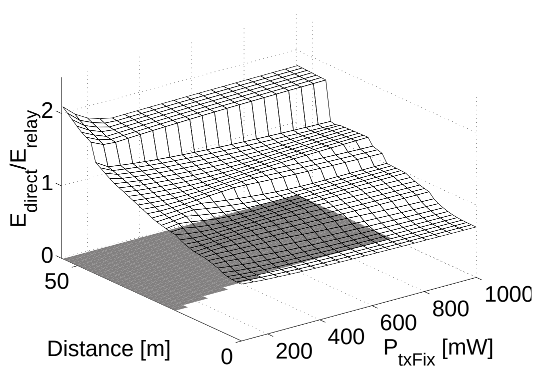

Figure 3.1: Ratio between E_{direct} and E_{relay}

Figure 3.2: Lifetime extension, a function of G

3.1 Energy efficient relay selection

Based on these calculations, a relay algorithm is simple to derive. The terminal with the lowest E_{relay} is used as a relay for a relayed terminal, if this energy is lower than E_{direct}. In addition, communication in HiperLAN/2 requires the distribution of the schedule from the AP, causing some small additional overhead [10] for the relaying case, and also causing an additional term in Equation (3.2).

Evaluating this algorithm in a scenario as described in Section 2 at a total load of about 3.8 MBit/s (two packets per frame and terminal)shows that energy efficiency increases (averaged over 55

[2] Parameters: $\alpha = 3.2$, target PER = 1%, $P_{\text{rx}} = 100$ mW, $a \approx 3.45$, two packets per frame. Gray area in the bottom plane indicates $E_{\text{direct}} > E_{\text{relay}}$.

different terminal placements) by only about 0.6 percent even though, in some placements, the gain is up to 10.8%. In fact, relaying is only effective in 16 out of 55 scenarios. The reason for this effect is the comparably large distance required for relaying to be beneficial—a condition that is only rarely met in scenarios with a maximum distance of about 50 m. Also, as all terminals have to be able to receive from the AP, larger distances are not feasible at the given α and target PER.

3.2 Relay selection considering battery capacity

Energy per bit is a system perspective, not immediately evident to the user. From a user's perspective, energy consumption is relevant as it reduces the time the terminal is functional. This "lifetime" of a terminal depends on the consumed energy, but also on the battery capacity. As some terminals have a large amount of battery capacity (e.g., terminals installed in cars), these could help out other users to improve the lifetime of their terminals.

Technically, this means that the relaying algorithm should not solely be based on the energy efficiency as such, but should also take into account the available battery capacity when choosing a relay terminal. More specifically: a terminal is considered to be a relay candidate for a far terminal if the following inequality holds:

$$\frac{C_{\mathrm{REL}}}{C_{\mathrm{FAR}}} > G \cdot \frac{E_{\mathrm{relay}}}{E_{\mathrm{direct}}} \tag{3.3}$$

where C_{FAR} and C_{REL} are the battery capacities of the far and the candidate relay terminal, respectively. G is a scaling factor and can be interpreted as the willingness of the algorithm to use non-optimally placed relay terminals when they have a large battery available. Among all candidates, the terminal with the smallest required energy to communicate with the far terminal is chosen as relay. This way, far terminals have a higher chance of using small transmission power and the load on relay terminals is spread according to their battery capacities.

To evaluate this algorithm, the smallest lifetime of all terminals in a cell using either direct communication or battery-aware relaying are compared. Figure 3.2 shows the lifetime extension achievable with different values of G where lifetimes are again averaged over 55 different scenarios and 9 MT's transmitting to the AP with a fixed data rate. Even though in this scenario, the energy efficiency can only be marginally improved, it is indeed possible to extend the lifetime of the network as a whole by more than 18% (the precise value depends on parameters like covered area, PER, and α).

Chapter 4

Optimizing capacity

The reduced distances made possible by relaying can improve the capacity of a wireless cell—here defined as the total amount of data that can be successfully handled by the access point per unit time—in two different ways: one is the reduced interference in neighboring cells (because of smaller transmission power), the other is the possibility of using faster modulations over these shorter distances.

Here, we are not focusing on the interference reduction resulting from relaying (this effect has been investigated in [11] and, for CDMA-based systems, in [1]), but here we are rather interested in the possibility to improve capacity by optimizing transmission schedules and modulations and by adding additional wireless resources. Relaying problems akin to this one have been studied in a number of publications (e.g., [5, 7, 8]), but the goals and system models are usually somewhat different: In particular, the combination of ad-hoc relaying as a means to amend a fixed network with a joint optimization of routing, modulation, and scheduling has—to our knowledge—not been studied before. In general, though, we corroborate the overall tendency of all papers that relaying can improve wireless capacity, but only up to a point.

4.1 System modes

Assuming a TDMA-like system, terminals like those in Figure 1.1 need to send their traffic in a given time slot. How to fairly divide the system resources (i.e, the time the AP can receive) between terminals—even in the direct case—gives raise to two different ways of contemplating the system. Fairness among terminals can be maintained by scheduling the communication such that either all terminals obtain an equal share of the total frame time (the “uniform slot size” scheme) or all terminals are allowed to send a constant amount of data in slots of varying length depending on their modulation and, basically, their distance from the access point (the “uniform traffic size” scheme); see Figure 4.1a. While the first scheme will result in a higher total goodput (as far terminals are given a relatively smaller weight), the second scheme corresponds better to user expectations. These two fairness schemes also extend to the relaying case: either all slots have the same length, or time slots are arranged such that all terminals receive the same effective goodput, no matter whether they are relayed or not.

In simple relaying, terminal A waits for an exclusive time slot to send data to B. A second time slot is then needed to communicate A’s data from B to the AP (where a faster modulation can be used in these time slots). However, C's sending to the AP and A's sending to B involve different entities and could—in principle—be scheduled concurrently. Doing so on the same frequency might be possible in large cells, but in order to avoid interference within a cell, another frequency could be used for the relaying step. Switching between two frequencies is indeed feasible as it is possible to have frequency switching times in ranges of few microseconds; also, no additional synchronization is necessary as switching happens within a frame. Figure 4.1b and c show the different fairness schemes for both one- and two-frequency relaying cases. Using an additional frequency might be inconvenient when this

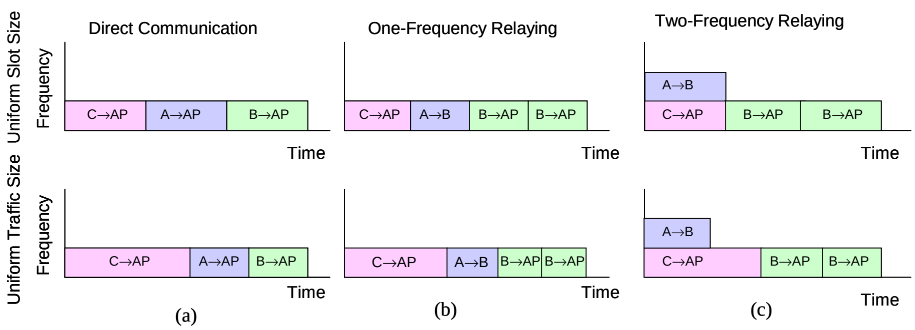

Figure 4.1: Time-freqency representation of the example scenario

frequency is used in neighboring cells. When there are sufficiently many frequencies available, this is not necessarily a problem: HiperLAN/2, e.g., specifies 19 frequencies, which should be sufficient for most reuse patterns. Moreover, the use of the second frequency can be confined to the communication of the AP with nearby terminals, which takes place at small transmission power values—from the outside, such a cell would indeed not disturb a frequency assignment too much.

4.2 Routing

The potential benefit of relaying depends on the data rates that can be realized between relayed terminal, relaying terminal, and access point. Hence, a joint optimization of routing and scheduling is necessary that decides which terminal to use for relaying and that selects modulation and transmission power. Currently, the algorithm optimizes only the uplink case and considers a single intermediate relay.

The effective data rate between two terminals can be determined based on their distance and a target packet error rate, allowing to compute, for each modulation, the required transmission power (based on results in [12]). Any modulation that requires more than a maximum allowable power or that does not match receiver sensitivity is ruled out, resulting in the optimum modulation for this pair of terminals; the transmission power is then adjusted to the smallest value that still meets the target PER for this modulation. This determines the modulation to be used in the scheduling decisions.

The routing then proceeds as follows: for a mobile terminal X, its data rate in the direct and relaying case is computed for all possible relaying terminals; the effective data rate for a relayed terminal is the minimum rate of both involved links. X selects that terminal as a relayer that results in the largest data rate for X and that exceeds the direct case data rate. This is a local routing decision only, global optimizations are left for further work; also, relaying terminals always communicate directly with the AP.

4.3 Scheduling mechanism

For both uniform slot size and uniform traffic size fairness schemes, relaying schedules are computed on the basis of the routing decisions. Let us first consider the uniform slot size case. In direct communication, all terminals are assigned an equal share of transmission time. In one-frequency relaying, additional time slots are needed for communication from the relayed to the relaying terminals, from the relaying terminals to the AP for their own traffic and from the relaying terminals to the AP for the relayed traffic. Each of these communications is assigned an equal share of time. For the two-frequency relaying, as communication from the relayed terminals to the relaying terminals can overlap with other communication, the total number of time slots depend on this overlap; and again, all slots are of uniform length.

In the case of uniform traffic size schedules, the optimal traffic size per slot depends on the modulations chosen in the routing phase. The schedule construction is in principle similar to the uniform slot size case, but here the slot length is individually varied for each terminal such that all terminals' goodput at the AP results in the same value, taking into account the need to relay traffic in two slots.

4.4 Results

For the model described in Section 2, the goodput achieved by capacity-oriented schedules for direct, one-frequency, and two-frequency relaying are averaged over 55 different placements of terminals. The resulting average total goodputs are shown, for ten entities and varying pathloss coefficient α, for both uniform slot size (Figure 4.2) and uniform traffic size schedules (Figure 4.3). For small α, all approaches are capable of fully utilizing the access point's maximum goodput. As α increases, the range of communication over which the 1% PER condition can be met at a fixed modulation and limited transmission power decreases. Hence, the observed goodput at the access point also goes down. When α exceeds 3.2, the 1% PER can no longer be met even with the slowest modulation, therefore, no values for larger α are shown.

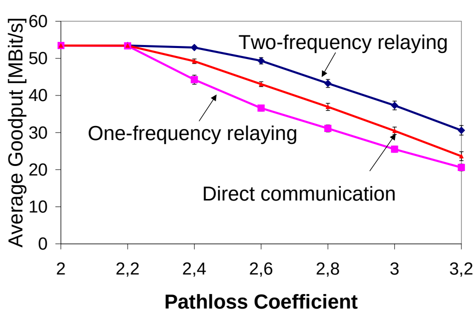

Figure 4.2: Average goodput as a function of α for uniform slot size fairness

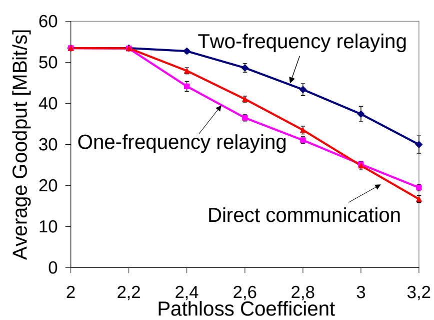

Figure 4.3: Average goodput as a function of α for uniform traffic size fairness

For both scheduling types, the goodput achieved by one-frequency relaying is (for this number of terminals) smaller than that of direct communication. This is due to the additional time slot requirement to send the relayed traffic to the access point. As α increases, one-frequency relaying approaches direct communication and actually surpasses it for uniform traffic size schedules (this tendency continues to larger α when the PER requirement is relaxed). In either case, two-frequency relaying considerably improves the total performance.

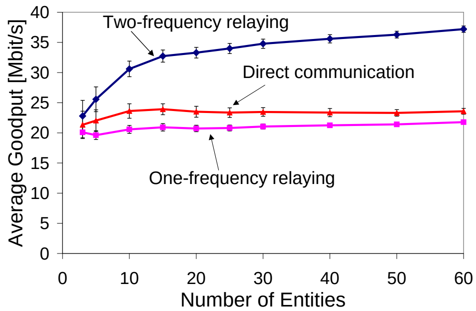

Figure 4.4: Average goodput as a function of entities for uniform slot size ($\alpha = 3.2$)

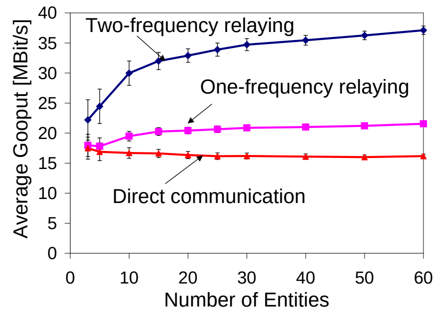

Figure 4.5: Average goodput as a function of entities for uniform traffic size ($\alpha = 3.2$)

This somewhat ambivalent result is partially due to the number of terminals used in this example. Figure 4.4 and 4.5 study the impact of varying the number of terminals. While in the uniform slot size case, one-frequency relaying still does not compare favorably with the direct case, uniform traffic size scheduling results in an improvement already for a modest number of terminals. Again, in both cases, two-frequency relaying results in considerable improvements, almost doubling the direct case's goodput in extreme examples. This dependency on the number of terminals is essentially a stochastic effect: the far terminals have a higher probability of finding a relay terminal, more often enabling faster modulations. Hence, relaying does generate more capacity when it is most sourly needed. Of course, relaying is not capable of improving the goodput per terminal as the AP is the bottleneck, in accordance with theoretical results.

The performance results for one-frequency relaying also show that the local heuristic used to determine relaying terminals needs to be extended for the uniform slot size scheduling case in order to also result in a *global* improvement of goodput. This is a challenging issue for future work.

Chapter 5

Electromagnetic field immission

As relaying reduces transmission power, it should also be a viable technique to lower electromagnetic field (EMF) immissions, which is currently often considered a health hazard. We examined the electromagnetic immission (received power) when using our energy- and capacity-optimized algorithms; moreover, we also developed a relaying algorithm that minimizes the radiated energy (power multiplied by time) while keeping the target PER.

While executing the simulations described in Sections 3 and 4, the EMF immission power is recorded by a regular grid of "sensors", placed 10 m apart. Relaying-based immission power is then compared to the direct case using two different metrics: First, the time-weighted average of the immitted power is calculated for all sensors (averaged over all topologies). Second, the sum of the time during which the immitted power exceeds a threshold is calculated (averaged over all topologies for each sensor).

Regarding the first metric, in Figure 5.1 the EMF immission ratio of the radiated-energy-optimized algorithm to the direct case is shown. The dark surface shows the region at which the immitted EMF is reduced (ratio < 1). This algorithm reduces the immitted power value by about 45%, averaged over all sensors and topologies. The curves for other algorithms (with respect to direct communication) are similar and they all do reduce the average immission power between 1.4 % (for the battery-optimizing algorithm) and 11 % (for the energy-optimizing algorithm). The loads used in these evaluations were the same as in the capacity and energy-efficiency evaluations, respectively.

The second metric's purpose is to reflect the current debate about the precise medical relevance of immitted power levels. As so far no safe threshold is established, we show the amount of time that a sensor is exposed to power levels exceeding an arbitrary threshold. To compare different algorithms, the resulting times are normalized to a single MAC frame length (2 ms). The results are shown in Figure 5.2; as an example, an average sensor is exposed to immitted power larger than 50 nW for 50 μs in every MAC frame when using the emission-optimized relaying.

One conclusion is that the characteristics of the curves are not changed by relaying, yet there is an improvement for all thresholds and algorithms. As an outcome we can improve capacity or energy efficiency while simultaneously reducing both average and peak immitted power. In most cases, immission power values at the edges are smaller (thus also reducing the inter-cell interference) and somewhat increased inside, near to the AP.

Note that these measurements are only for uplink data transmissions. In the downlink case, we expect a decrease in immitted power since the AP itself transmits for a longer time such that power reduction should have a considerable effect.

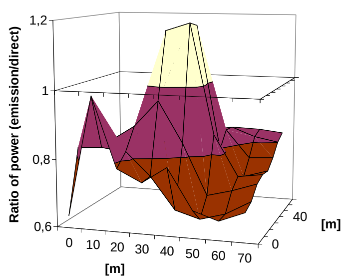

Figure 5.1: EMF Ratio of emission-optimised relaying and direct mode operation for each sensor position, averaged over scenarios

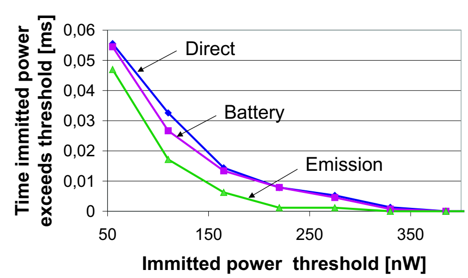

Figure 5.2: Amount of time during a MAC frame for which the threshold is exceeded, averaged over all sensors and scenarios

Chapter 6

Conclusions and outlook

The concepts and algorithms presented in this paper show that relaying is a viable means to improve the operations of an infrastructure-based wireless communication system: it is possible to extend the lifetime of the network, considerably increase the cell capacity or reduce the electromagnetic immission. Moreover, the algorithms are practical as they are not computationally intensive, they can be implemented as online algorithms iteratively improving intermediate solutions, and are based on information (in particular, channel gain between terminals) that can be provided by real systems with acceptable overhead (e.g., HiperLAN/2's radio map). We are currently implementing (together with partners from the HyperNET/IBMS2 project) relaying in a real HiperLAN/2 testbed and will study its performance.

A number of questions remain open: foremost, the impact of mobility should be investigated. Preliminary studies show that the algorithm performance is robust up to medium speed and that the update interval for channel gain measurements is acceptable. Also, the possibilities of relaying-based immission reduction warrant additional investigations. Furthermore, the effects of interference reduction and usage of multiple frequencies will be studied in a multi-cell context. An issue of larger magnitude is the integration of directed (smart) antennas, which should bring additional benefits to all aspects of the relaying problem.

Bibliography

[1] M. Bronzel, W. Rave, P. Herhold, and G. Fettweis. Interference reduction in single hop relay networks. In *Proc. of 11th Virginia Tech Symposium on Wireless Personal Communications (to appear)*, Blacksburg, VA, June 2001.

[2] B. Burns and J.-P. Ebert. Power Consumption, Throughput and Packet Error Measurements of an IEEE 802.11 WLAN Interface. Technical Report TKN-01-007, Telecommunication Networks Group, Technische Universität Berlin, August 2001.

[3] T. A. ElBatt, S. V. Krishnamurthy, D. Connors, and S. Dao. Power management for throughput enhancement in wireless ad-hoc networks. In *ICC 2000*, New Orleans, LA, June 2000. IEEE.

[4] L. M. Feeney and M. Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. In *INFOCOM*, Anchorage, AK, US, April 2001. IEEE.

[5] M. Gastpar and M. Vetterli. On the capacity of wireless networks: The relay case. In *Proc. INFOCOM*, New York, NY, June 2002.

[6] J. Gomez, A. T. Campbell, M. Naghshineh, and C. Bisdikian. Conserving transmission power in wireless ad hoc networks. In *9th International Conference on Network Protocols*, Riverside, California, November 2001. IEEE.

[7] M. Grossglauser and D. Tse. Mobility increases the capacity of wireless adhoc networks. In *Proc. Infocom*, 2001.

[8] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. on Information Theory*, November 1998. `http://black.csl.uiuc.edu/~prkumar/ps_files/capacity_final.ps`.

[9] D. Hollos and H. Karl. A HiperLAN/2 simulation model in OMNeT++. In *Proc. 2nd Intl. OMNeT++ Workshop*, pages 61–70, Berlin, Germany, January 2002.

[10] D. Hollos and H. Karl. A protocol extension to hiperlan/2 to support single-relay networks. In *Proc. of 1st German Workshop on Ad-Hoc Networks*, pages 91–108, Ulm, Germany, March 2002. `http://www-tkn.ee.tu-berlin.de/publications/papers/frame2.pdf`.

[11] H. Karl and S. Mengesha. Analyzing capacity improvements in wireless networks by relaying. In *Proc. IEEE Intl. Conf. Wireless LANs and Home Networks*, pages 339–348, Singapore, December 2001.

[12] J. Khun-Jush, G. Malmgren, P. Schramm, and J. Torsner. HIPERLAN type 2 for broadband wireless communication. *Ericsson Review*, 2:108–119, 2000. `http://www.ericsson.com/review/2000_02/files/2000026.pdf`.

[13] J. Peetz. HiperLAN/2 multihop ad hoc communication by multiple-frequency forwarding. In *Proc. of Vehicular Technology Conference (VTC) Spring 2001*, Rhodes, Greece, May 2001.

[14] Y. Xu, J. Heidemann, and D. Estrin. Geography-informed energy conservation for ad hoc routing. In *Seventh Annual ACM/IEEE International Conference on Mobile Computing and Networking*. ACM, July 2001.