TKN **Telecommunication Networks Group**

Technical University Berlin

Telecommunication Networks Group

Virtual Utilization and VoIP Capacity of WLANs Supporting a Mix of Data Rates

Sven Wiethoelter

wiethoel@tkn.tu-berlin.de

Berlin, September 2005

TKN Technical Report TKN-05-04

TKN Technical Reports Series

Editor: Prof. Dr.-Ing. Adam Wolisz

Abstract

Today's WLANs offer high data rates of up to 11 and 54 Mbps for IEEE 802.11b respective 802.11a/g making this technology very attractive for real-time applications like VoIP calls. However, only a certain part of communication can take place at high data rates. Additionally, a certain amount of overhead is introduced by the CSMA/CA-based medium access scheme. Thus, challenging questions remain how the load situation of a WLAN cell can be evaluated adequately and which costs arise if some users transmit only at low rates.

In this work, we motivate a "virtual" utilization metric for the evaluation of the current situation in a WLAN cell and investigate simulatively the influence of low-rate transmissions in a WLAN network used for pure VoIP traffic .

The virtual utilization is based on the physical and virtual occupancy duration for each transmission. Since the occupancy duration is measured at the AP, which is the bottleneck of a WLAN cell, we are able to identify the resources that every user consumes for his transmissions—including the effects of transmission rates as well as technology dependent overheads.

For the VoIP simulation studies, our metric is the number of calls that is servable at acceptable quality for every single user. "Acceptable" thereby means that certain VoIP QoS constraints are not violated. For each single data rate, the VoIP capacity, i.e., the maximum number of calls, is determined at first. Afterwards, the influence of a single low-rate user is investigated and set in relation to the capacity results, which leads to the identification of costs invoked by the low-rate transmissions.

Contents

Chapter 1

Introduction

In today's WLAN, users mostly have nomadic mobility degree, i.e., they appear in the cell and remain at a constant position during their connectivity. Since the wireless channel is error-prone, the situation of a user can nevertheless change due to multi-path fading effects although he keeps a constant position. In the very near future, dual phones, which are mobile phones with a cellular and a WLAN interface, will emerge, thus nomadic users as well as users with pedestrian mobility will reside inside a WLAN network. Since mobile users may arrive at or leave a cell during an on-going application session like a VoIP call, a large challenge is to enable seamless handovers not only among homogeneous WLANs but also among heterogeneous WLAN /UMTS networks.

In such a scenario, the question arises when a handover should be performed. A simple approach could decide e.g. dependent on the Received Signal Strength (RSS). A handover would be performed, if the average power level of a received signal falls below a certain threshold. However, WLAN devices which are capable to transmit at several high-rate modulation schemes try to adapt their transmission rate to the current situation, i.e., a mobile user moving away from AP throttles down its data rate further and further to gain error-free transmissions. This behavior continues until the received signal at the mobile terminal undergoes a threshold indicating that a HO should take place.

A threshold-based handover has the disadvantage that it can affect other terminals since a HO decision is made when the lowest transmission rate cannot be used for error-free transmissions anymore. Transmissions at lower rate lead to increased transmission times and may consume resources which are badly needed to serve other users. Even if the low-rate mobile can be served such that the perceived quality of its application is sufficient, it can influence other users' perceived QoS. In an extreme case, transmissions at a degraded rate result in service interruptions for others.

This study investigates the influence of different transmission rates in an IEEE 802.11b WLAN for pure VoIP traffic. The metric under study is the number of calls that is possible in combination with acceptable quality for all VoIP users. Firstly, we consider the capacity of WLAN for VoIP at each single data rate. Afterwards, we identify the costs invoked by one user transmitting at a lower rate.

Two different codecs—ITU-T's G711 and G729—with packetizations of 10 to 50 ms are considered in our simulation runs, which have been repeated independently forty times to gain significant results. For our investigations we assume that no fading effects occur and

stations adapt their transmission rates perfectly such that the BER remains always below 10^{-5}. Therefore transmission failures due to bit errors are negligible. These assumptions are made since the minimum costs are of interest, i.e., whether high degradations as a result of lower rates occur even for a perfect system.

In order to allow decisions for handover and admission control algorithms, a criterion is presented which takes into account the overhead as well as the modulation scheme of each transmission. It is based on the physical and virtual occupancy time of the AP that is necessary for each up and downlink transmission.

The rest of this paper is organized as follows. Chapter 2 describes the overhead of transmissions involved by the IEEE 802.11 protocol and rate adaptation mechanisms. Chapter 3 shows that the WLAN overhead greatly reduces the number of possible calls and presents previous work identifying the number of calls in a WLAN cell. Chapter 4 introduces in detail the evaluation criterion and the choice of the perceived quality criteria. The VoIP Network stack, the traffic model, the perceived quality constraints for VoIP, the scenario and the simulation model are discussed in Chapter 5. Chapter 6 and Chapter 7 present the VoIP simulation results of the G711 respective G729 codec. The results of the virtual utilization and potential application fields are discussed in Chapter 8. Finally the conclusions are drawn in Chapter 9.

Chapter 2

WLAN Overhead

This chapter first describes the medium access scheme of IEEE 802.11 and the Physical Layer enhancement IEEE 802.11b by focusing on the overhead that is introduced for each transmission. Since no rate adaptation mechanisms for high-rate enhancements have been specified by the IEEE, published approaches are discussed to illustrate their overheads.

2.1 IEEE 802.11 MAC

IEEE 802.11 [2] defines two medium access functions: the Point Coordination Function (PCF) and the Distributed Coordination Function (DCF). While the DCF is responsible for asynchronous data services, the PCF offers time-bounded services. The PCF is used in the Contention Free Period (CFP), while DCF handles medium access in the Contention Period (CP). One CFP and one CP are combined to a superframe, which repeats periodically.

2.1.1 DCF

In DCF, sending stations stand in contention to each other. Due to Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA), a station has to sense the channel before being allowed to send a frame. The station may send if the medium has been idle for at least Distributed Interframe Space (DIFS). In case of a busy medium, the station starts a random backoff procedure—the binary exponential backoff. If the medium is idle again at least for DIFS, the station decrements the backoff time until the medium gets busy again or the residual backoff time is equal to zero. In the latter case, the station is allowed to send out its data packet immediately. If the receiver got the packet without any errors, it sends out an Acknowledgment (ACK) after Short Interframe Space (SIFS). Since the duration of SIFS is smaller than DIFS, the immediate ACK transmission has priority over all other attempts to access the channel.

The sender starts a random backoff procedure after every transmission. In case of a successful acknowledged transmission, the procedure starts after the received ACK (denoted as postbackoff). Otherwise the backoff is invoked after the expiration of the ACK timeout interval. The sender assumes in this case that a collision happened. A collision occurs if two (or more) stations have detected the medium as idle for DIFS, both are allowed to send and both start their transmissions immediately. The contention window, out of which the

random backoff period is drawn, doubles for each unsuccessful transmission attempt to avoid repeating collisions among stations.

For each unicast transmission, the IEEE 802.11 MAC introduces a certain amount of overhead, i.e., MAC headers, interframe spaces, immediate ACK, as well as backoff(s) as described above. A MAC header of 34 Bytes adds to each MAC Service Data Unit (MSDU). Due to the CSMA/CA scheme, a station has to sense the channel for DIFS prior transmission and has to wait for SIFS before the receiver is allowed to start sending its ACK. After every successful transmission as well as for each missing ACK, a station performs a backoff which further increases the overhead.

2.1.2 PCF

The PCF can only be used in an infrastructure-based network because it requires a central coordinator—the Point Coordinator (PC). Usually the PC is installed on the AP and manages the access to the medium in the CFP by polling stations sequentially. This results in a high and potential unnecessary overhead if the polled station has no packets ready to transmit. The PCF comes up with a higher complexity than the DCF. It is mostly implemented in currently installed infra-structured BSS but it is not used very often due to the lack of an optimized scheduling / polling.

2.2 IEEE 802.11b PHY

The IEEE 802.11b [3] supplement enhances the 802.11 standard by two high rate capabilities. The Complementary Code Keying (CCK) modulation scheme provides data rates of 5.5 and 11 Mbps. However, not the whole PLCP Protocol Data Unit (PPDU) is transmitted at these high data rates. Since downward compatibility to legacy 802.11 devices has always been one design goal, the long PHY headers of 802.11b consisting of the same Physical Layer Convergence Protocol (PLCP) preamble and the PLCP header are transmitted at the basic rate of 1 Mbps with DBPSK, too. 802.11b specifies additionally a short PPDU format, which can be used in pure high-rate WLANs only. Since this short format is optional, we do not further consider it in our studies. Like in legacy 802.11, the higher-rate modulations can be used only for the PLCP Service Data Unit (PSDU) part. This leads to a non-negligible constant overhead of 192 μs for each frame transmission.

2.3 Rate Adaptation Mechanisms

Today, several enhancements exist to the first WLAN PHYs specified in the IEEE 802.11-1999 standard [2]. Like in IEEE 802.11b, a set of data rates is available from which each station chooses one rate for the transmission of the MAC Protocol Data Unit (MPDU). However, high-rate modulations are more susceptible to impairments like multi-path propagation, noise, and path loss. Thus, a station should adapt its transmission rate to the current situation at its receiver. The sender should find the optimal rate in order to avoid on the one hand a waste of resources due to an unnecessary robust modulation while on the other hand a falsely chosen higher rate results potentially in more retransmissions. No rate adaptation

TKN-05-04   Page 6

mechanisms for WLAN have been specified by the IEEE so far, so each manufacturer chooses its on proprietary solution. In the following, published schemes are discussed for an overview.

Kamerman et al. [14] described the rate adaption scheme that is used in WaveLAN-II devices. The Automatic Rate Fallback (ARF) scheme works as follows. If several successful transmissions are followed by two consecutive transmission failures (indicated by missed ACK), the data rate is degraded down to the next lower rate for all subsequent (re)transmissions. With the change to the fallback rate, the station starts a timer. If either the timer expires or ten consecutive successful transmissions (10 error-free ACKs) occurred, the data rate is increased again to the next higher level. If the following transmission fails, the fallback rate is used again and the timer / ACK-count procedure restarts.

Two problems of ARF were identified by Lacage et al. [17]: This method performs non-optimal if channel conditions either change quickly or remain constant over a long period of time. If channel conditions vary fast, ARF never adapts to the current situation. In the latter case, since ARF tries to upgrade its rate after every tenth successful transmission, a higher overhead occurs due to an increased number of retransmissions. Therefore, the authors propose Adaptive ARF (AARF) which enhances ARF by introducing an exponential binary increment of the successfully received ACKS if the first transmission fails after switching to a higher rate. Via simulations, improvements over ARF were identified while the authors showed in a testbed with a modified MADWiFi [1] driver that their solution outperforms the legacy MADWifi rate adaptation which is based on packet loss rates.

Since the optimal transmission rate depends on the channel state at the receiver, Holland et al. [10] proposed Receiver-Based Auto Rate (RBAR) which exploits this fact by combining rate adaptation with an RTS/CTS handshake. Dependent on the RTS (e.g., its received signal strength), the receiver chooses the transmission rate for the data packet and includes this information into its CTS. RBAR does not only require an RTS/CTS handshake but also needs special header fields thus changing the legacy IEEE 802.11 MAC and PHY headers.

Another approach [20] is based on RSS measurements too, but additionally assumes symmetric channel conditions. Pavon et al. proposed an algorithm that averages RSS information of previously received frames from the AP. Beside the RSS measurements, it includes an adjustment of rate thresholds, which have been investigated simulatively, dependent on the number of retransmissions.

---

[1]Multiband Atheros Driver for WiFi

Chapter 3

Interaction of WLAN and VoIP

The first part of this chapter shows a very simple calculation of the maximum number of VoIP calls if only the high data-rate and the netto bandwidth of a voice call is considered. A theoretical investigation follows which shows the impact of framing and MAC / PHY overhead. Lastly, previous work is presented which focusses on the number of possible VoIP calls in a WLAN environment.

3.1 Misleading Marketing Slogans

Today's marketing sloglans for WLAN network interface cards might give customers the impression that the technology is always able to gain high data rates of up to 11 or 54 Mbps for IEEE 802.11b respective 802.11a/g. Using VoIP over WLAN, customers could be tempted to calculate the number of possible voice calls in a WLAN cell as follows: Applying, e.g., the G711-audio codec, results in an audio-data rate of 64 kbps. Due to its bi-directional nature, one call produces an audio rate of 128 kbps. Neglecting the overhead of lower layers and taking into account WLAN's data rates, a basic customer would assume that 85 VoIP calls are possible with 11 Mbps [8] and 421 calls with 54 Mbps.

The above estimations disregards the important aspects that WLANs come up with a high amount of overhead due to their medium access scheme and the fact that the PHY transmits only a certain part of a frame with the high data rate.

3.2 VoIP Capacity with Overheads

A short numerical and more realistic investigation of the capacity for VoIP over 802.11b WLAN is applied in the following. Again, we consider G711-coded speech with a bandwidth of 64 kbps. The audio stream is separated in fragments with a length of 20 ms, so that each fragment consists of 160 Bytes.

The first part of the overhead is introduced due to framing at different layers. To each fragment, higher layers add 40 Bytes for UDP/RTP/IP headers. Thus the frame has already a size of 200 Bytes when the packet arrives at the 802.11 MAC. MAC and PHY themselves add headers of 34 Byte respective 24 Bytes, thus the VoIP frame on air has a size of 258 Bytes. As explained in Chapter 2.2, the PHY only transmits the PSDU itself (i.e., 234 Bytes) at

Table 3.1: Overheads for each VoIP packet transmission, G711-coded speech, 20 ms packetization

audio payload	160 Bytes
packetization	20ms
RTP + UDP + IP header	40 Bytes
MAC header for DATA	34 Bytes
MAC header for ACKs	14 Bytes
PLCP preamble and header	24 Bytes
slot duration	20 μs
SIFS	10 μs
DIFS	50 μs

the high data rate, while the preceding PLCP preamble and header are transmitted with the basic rate.

Due to the medium access scheme, interframe spaces, immediate ACK, and (post)backoff increase the overhead for the transmission of every VoIP packet. For an easy overview, all relevant parameters for a transmission are listed in Table 3.1.

Since stations can decrement their backoff simultaneously, we neglect any retransmissions and postbackoff and compute simply the occupancy time consisting of inter-frame spaces as well as the transmission duration for VoIP packet and immediate ACK for each transmission. The medium is completely occupied if only 13 to 14 calls are served in an IEEE 802.11b network. The number of calls will further decrease by taking into account the post-backoff and transmission failures like collisions and bit errors.

The authors of [8] applied a similar investigation by taking into account additionally a mean occupancy time for collisions and backoffs which has been gained via analysis and simulations. With this, the autors calculated that 12 VoIP calls in total can be served with G711-coded speech and packet-generation intervals of 20 ms in a 802.11b network.

3.3 Number of VoIP calls

Koepsel et al. [15] compared the usability of PCF and DCF for time-bounded applications in simulative studies. Among others, the authors considered PCM-coded VoIP traffic modeled by a two-state Markov model, i.e., silence suppression has been enabled. With a packetization of 20 ms, a maximum one-way delay of 250 ms and an upper bound of 5 percent packet loss, DCF supports 12 one-way audio flows (i.e., 6 calls), while PCF is able to service 15 flows in a 2 Mbps WLAN. The authors concluded that the throughput degradation due to unsuccessful polling attempts in PCF mode could be reduced by optimized polling structures, i.e., the usage of fine-grained scheduling algorithms as well as polling lists.

Today's WLANs mostly use the DCF to regulate the medium access since PCF lacks optimized scheduling / polling mechanisms as proposed by [15]. Therefore, we focus in the following on previous work which identified the maximum number of VoIP calls in WLANs with DCF medium access method only.

In [8], Garg et al. considered the time that a voice call utilizes a WLAN cell by taking into account not only the transmitted VoIP packets but also the resulting overhead like ACKs, interframe spaces and backoff durations. For 802.11a/b WLANs, the authors calculated the maximum number of VoIP calls for 3 different codecs: G711, G729 and G723.1. All codecs were investigated without silence suppression and in combination with various packetizations ranging from 10 to 100 ms. In our work, the 802.11b results at a data rate of 11 Mbps of [8] serve as a basis for verification of our simulation model. In a follow-up publication, Garg et al. [7] applied throughput measurements for UDP and VoIP traffic in an 802.11b testbed. With G711-coded speech and a packetization of 10 ms, the WLAN network was able to serve 6 calls. The authors observed packet loss, delay and its jitter to decide for acceptable call quality.

Elaoud et al. [5] applied a traffic generator with a 4-state Markov model for their measurements with G711-coded voice and 10 ms packetization. Their metrics packet loss and one-way delay had boundaries of 2 percent and 200 ms respectively. In the testbed, not only a pure WLAN but also a LAN and WAN environment was emulated additionally. Considering packet loss only, 10 VoIP calls were able in the 802.11b network, while 6 calls and 4 calls could be served in combination with the delay invoked by LAN and WAN environment respectively.

Voice capacity of IEEE 802.11a/b/g WLANs has been investigated analytically and simulatively by Medepalli et al. in [18]. For G711-coded speech with silence suppression (modeled by a 4-state Markov model) and packetizations ranging from 10 to 50 ms, the maximum number of calls have been identified for single 802.11a/b/g as well as for combined 802.11b/g networks. In simulations, the metrics packet loss and one-way delay have been chosen with maximum values of 2 percent and 200 ms respectively. Compared to a 802.11b network, 802.11a/g are able to serve approximately 5 times more calls. Since specific means like slot-time adaptation as well as RTS-to-self or RTS/CTS protection are needed in a combined 802.11b/g network, the number of calls drops significantly in this environment. In a follow-up publication [19], Medepalli et al. considered the voice capacity in 802.11a/b WLANs in case of channel errors and if certain users incorporate a different data rate. The analysis of [18] was based on the expected time that the network is physically and virtually occupied (similar to [8]). In [19], the model was extended by including channel errors and users at different rates. For PCM speech with a packetization of 30ms, the numbers of highest-rate vs. lowest-rate users were gained for 802.11a/b: One lowest-rate user (802.11b: 1Mbps, 802.11a: 6Mbps) reduced the number of highest-rate users (802.11b: 11Mbps, 802.11a: 54Mbps) by 3.

Chapter 4

Virtual Utilization—Principle

A proper metric that reflects the current situation of a WLAN cell has to cover different transmission rates and overheads of every transmission. Thus, an approach which is based on the occupancy time for each transmission is very attractive. In [6], such an approach has been already published. The authors apply a measurement which they denote as "Channel Utilization Estimate (CUE)". It is defined as "the fraction of time per time unit needed to transmit the flow over the network". CUEs of all flows in the network are simply cumulated thus specifying the overall fraction of time that the medium is busy—which gives the overall utilization.

CUEs are either calculated on the basis of simply counting transmitted packets and their overhead for every flow or by taking into account mean values of backoff parameters and mean collision rates. These parameters are then used to estimate the CUE values respective to define upper bounds for the maximum, accumulated CUE over all flows. It is not specified in [6] precisely, how the backoff scheme is taken into account for the CUE computation. A problem arises if multiple stations are in a backoff process simultaneously. In case of an idle medium, they will decrement their backoff counter simultaneously until one station has completed its backoff period and starts its transmission. If the parallel backoff decrement would be taken into account for each station transmitting on the uplink, it would falsify the occupancy-time computation. In other words, the influence of the backoff would be overestimated.

4.1 Utilization of the WLAN network and Load of the AP

In infra-structured WLAN networks, one AP serves all wireless stations, i.e., it has to contend for channel access prior every transmission in the downlink direction. In legacy IEEE 802.11, the MAC parameter sets for AP as well as wireless stations are equal which results in a fair medium access with respect to the number of transmitted packets. As a result, the AP becomes the bottleneck in the downlink direction since it has to serve all other wireless clients. Therefore we consider the load of the AP by taking into account

- the duration of all transmissions in up- as well as downlink direction,

- plus the backoff duration for each downlink transmission at the AP.

The above described durations are simply cumulated and averaged over a certain time interval Δt which leads to the virtual utilization U_{AP} of the AP (Equation 4.1).

$$U_{AP} = U_{uplink} + U_{downlink} + U_{backoff} \tag{4.1}$$

, with $\ U_{uplink} = \dfrac{t_{busy_{uplink}}}{\Delta t}, \ U_{downlink} = \dfrac{t_{busy_{downlink}}}{\Delta t}, \text{and} \ U_{backoff} = \dfrac{t_{backoff}}{\Delta t}$

The utilizations U_{uplink} and $U_{downlink}$ contain the frame durations (t_{data}, t_{ACK}) plus all Inter-Frame Spaces (IFSs) (t_{SIFS}, t_{DIFS}) which are needed for the transmissions in up- and downlink directions, while $U_{backoff}$ only consists of the slot durations that the AP chose for its backoff during Δt.

U_{AP} reflects the busy time of the AP over a certain period Δt. The choice of Δt, i.e., the interval over which the virtual utilization is computed, is discussed further below in Section 4.2.

4.2 Averaging Interval

Averaging the busy time of the AP over a certain duration is used to eliminate factors which should not have a direct impact. Therefore, the averaging interval for the virtual utilization may depend on two issues:

1. slow-fading characteristics of the wireless channel, and

2. the generation rate with which the traffic is injected into the wireless network.

To be able to eliminate the effects of slow-fading, the coherence time is considered in the following. Assuming slowest movements in the considered environment of around 2 km/h, it is of interest for which mean duration the wireless channel remains approximately invariant. The coherence time is a statistical measure which describes the duration in which the channel impulse response is mainly invariant [22]. For this duration, the probability is high that the channel impulse responses do not vary greatly. The coherence time T_C as defined for "modern digital communications" in [22] is given in Equation 4.2,

$$T_C = \sqrt{\frac{9}{16\pi f_m^2}} = \frac{0.423}{f_m} \tag{4.2}$$

with f_m being the maximum Doppler shift

$$f_m = \frac{V}{\lambda} = \frac{V f_c}{c}$$

Since slowest movements will result in largest coherence times, a slow fading channel is considered. Slowest movements of around 2 $km/h = 0.556\ m/s$ in the environment are used to determine the coherence time (with center frequency $f_c = 2.4\ GHz$, velocity of light $c = 3\ 10^8 m/s$) which is in this case approximately equal to 95 ms.

Another effect which should be neglected are changes in the virtual utilization due the generation rate of data packets. For example, the time gap between two VoIP packets depends

on the VoIP codec and its packetization. Typical packetization values are between 10 and 50 ms.

As a result, the basis of the averaging interval will be chosen in the range of hundreds of milliseconds, thus eliminating changes in the virtual utilization due to slow-fading effects as well packetization.

4.3 Perceived Quality

Not only the utilization of a network but also the quality that the user of a specific service perceives needs to be taken into account. This is necessary to ensure a certain level of QoS.

Specific models which evaluate the perceived quality of speech transmissions exist already: ITU's PESQ model compares an original and a degraded speech sample, the resulting MOS value describes the residual quality of the degraded sample. PESQ does not take into account impairments like end-to-end delay, echoes, side tone, and loudness [9]. Therefore ITU-T specified the E-Model [13] which considers impairments that occur in typical telephone call conversations.

The disadvantage of both models is that they are not applicable in real-time due to the resources needed for the calculation, thus making a live evaluation of the perceived quality difficult if not even impossible.

To enable a time-aligned evaluation of the perceived quality, it seems to be most advantageous to consider at first goodput as well as the delay and its jitter for each flow. Defining maximum upper bounds that are service specific will enable a perceived quality evaluation.

Chapter 5

VoIP Simulation Model

5.1   VoIP Network Stack

For each call, a constant bit rate flow is set up in the up- as well as in the downlink direction. The voice stream will be encoded according either to ITU-T's G711 or G729 codec. G711 uses uncompressed PCM-coded speech that results in a bandwidth of 64kbps. Contrary, G729 compresses the voice signals which leads to a reduced bandwidth of 8kbps. For both codecs, VoIP packets usually include speech with a duration of multiples of ten milliseconds. Most commonly used values for the so-called packetization range from 10 to 50ms, whereby every 10ms result in 80 Byte voice data for G711 and only 10 Bytes for G729.

In order to meet the high requirements of voice transmissions, RTP provides timestamps as well as sequence numbers allowing to detect late transmissions or packet losses as well as to compensate packet reordering. On the transport layer, UDP enables a connection-less and unreliable service which is most beneficial for time-bounded applications like VoIP. RTP adds an overhead of 12 Bytes, while UDP enlarges the voice packet by 8 Bytes. Finally, IP adds additional 20 Bytes on the Network Layer. In total, this results in 40 Bytes overhead for each VoIP packet not taking MAC and PHY headers into account.

5.2   Metric and VoIP QoS constraints

The metric for our simulations is the total number of VoIP calls at acceptable quality. "Acceptable" thereby means that the VoIP QoS constraints are not violated. The rest of this section describes the choice of these constraints.

VoIP traffic is highly sensitive to delay, its jitter and packet losses. Assuming that echo cancellation techniques are applied[1], the maximum acceptable one-way delay is in the area of 200 [12] to 250ms [11]. If the delay exceeds this boundary, talkers would overlap each other thus resulting in a poor call quality. Actually, VoIP delay splits into several components. Despite networking delay, some latency is introduced by processing of VoIP samples, by decoders, and by playout buffers, which reduce the effect of jitter. Assuming that these additional latencies will not cross 50 to 100 ms, we limit the overall one-way networking delay to 150 ms (similar to [16]). If the latency of a VoIP packet exceeds this limit, we

---

[1]without this, the maximum allowable round-trip delay drops down to 50 ms [11, 16]

TKN-05-04   Page 14

assume that it results in such a large jitter that its content cannot be played out in time. As a result, these late packets are dropped.

Despite delay, loss of packets degrades the perceived quality of a VoIP call. According to [12], a one percent loss already leads to significant degradation. With Packet Loss Concealment (PLC) schemes this can be relaxed up to 5 percent for G711 [12]. Due to higher compression, other codecs are more susceptible. As a result, the maximum loss has been set to 5 percent in this work (similar to [16]).

If 5 percent of the VoIP packets are lost, i.e., they have been dropped or they arrive with a delay larger than 150 ms, the perceived quality is assumed to be lousy. The interval over which this criterion is evaluated has been set to 4 seconds. We have chosen this value in the order of seconds since the duration of lousy quality must be not only accustically perceivable for human-beings but also so large that the quality degradation is not acceptable anymore.

Table 5.1 lists all VoIP constraints for a short overview.

Table 5.1: VoIP QoS constraints

max. one-way delay	150 ms
max. lost packets	5 %
averaging interval	4 s

5.3 Scenario

The scenario comprises a wired and a wireless part. The latter one consists of one IEEE 802.11b Basic Service Set (BSS), i.e., one Access Point (AP) plus multiple wireless stations. The wired part consists of wired access networks, the backbone as well as wired clients which have VoIP calls with the wireless users. The behavior of the wired part is modeled with a constant packet latency of 100 ms.

The wireless part of the scenario under study is displayed in Figure 5.1. There, different modulation schemes are considered dependent on the distance of each wireless station to the AP. The area around the AP is divided into "modulation-scheme regions", whereby the modulation and therefore also the data rate within one region is modeled constantly. In Figure 5.1, these regions are displayed by circles around the AP.

In this investigation, all terminals have static positions only. The number of Stationary Terminals (STs) will be increased until the critical number of voice calls is reached, i.e., VoIP calls are dropped due to violation of the QoS constraints. Thereby, the maximum number of possible VoIP calls is not only gained at every single data rate, but also in situations where some terminals use a lower data rate.

The simulative study takes place by considering different transmission rates due to path loss. Each data rate is chosen such that a BER smaller than 10^{-5} is gained. Switching to the optimal rate is assumed to be perfect in a first step. Rate adaptation mechanisms like ARF, RBAR, etc. and packet losses due to bit errors are not modeled because the minimum costs of a VoIP call at degraded data rate are of interest for this work.
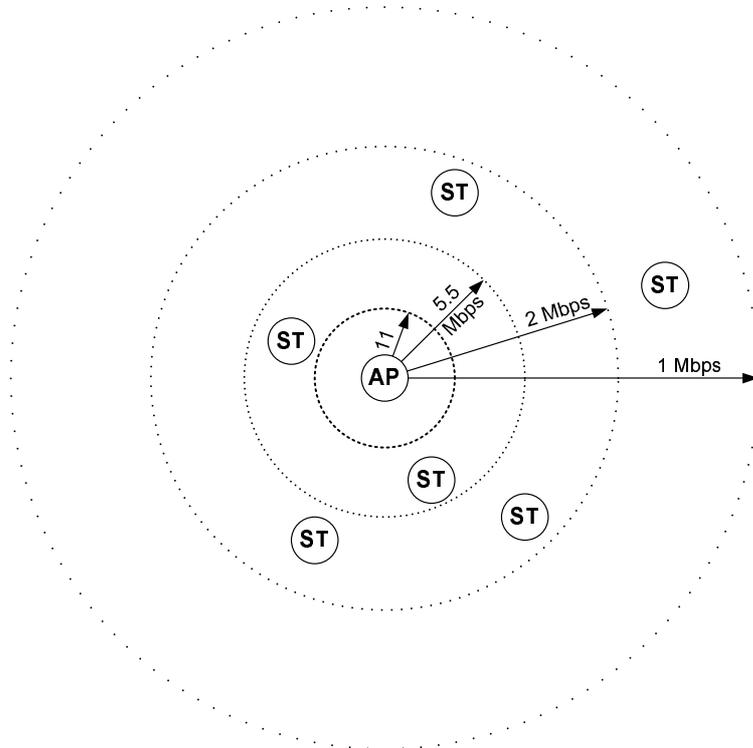
Figure 5.1: first scenario, different "modulation zones"

5.4 Modulation-scheme Regions

IEEE 802.11b uses four different modulation schemes, i.e., DBPSK (1.0 Mbps), DQPSK
(2.0 Mbps), 5.5 Mbps CCK , and 11.0 Mbps CCK. Modulation schemes are chosen in this
work such that the BER is below 10^{-5}. To allow a simplified determination dependent on
the BER, the 16-QAM is considered for the 5.5 Mbps mode while 256-QAM is used to model
the 11.0 Mbps modulation scheme (similar to [4]).

Considering an AWGN channel, the BER of DBPSK can be calculated by [21, p. 274]

$$BER = \frac{1}{2}e^{-\frac{E_b}{N_0}} \text{ , so that } \frac{E_b}{N_0} \geq -ln(2\,BER)\,. \tag{5.1}$$

Since the calculation of the relation E_b/N_0 is much more complicated for the other modula-
tions, they are gained either numerically via MatLab (for 16- and 256-QAM) or by metering
BER curves (for DQPSK) of [21, p. 275] . Table 5.2 lists E_b/N_0 for all four modulation
schemes. In this work, we consider a free-space propagation model, i.e., sender and receiver
have a direct line-of-sight path for the transmitted signal. The signal is only attenuated by
the path loss, so that the received power level can be calculated by [22]:

$$P_r(d) = \frac{P_t G_t G_r \lambda^2}{(4\pi)^2 d^2 L} \tag{5.2}$$

Assuming an AWGN channel as well as the parameters listed in Table 5.3, the distance
between sender and receiver has been calculated for all four modulation schemes such that

Table 5.2: Modulation-Scheme Regions

Modulation	$\frac{E_b}{N_0}$ [dB]	max. distance [m]
DBPSK	10.3	478
DQPSK	12.0	279
16-QAM	13.4	168
64-QAM	22.5	41

the power level at the receiver never falls below its sensitivity threshold. The maximum allowable distance for each modulation scheme is displayed in Table 5.2.

Table 5.3: Parameters for Modulation-Scheme Regions

transmit power	100 mW
receiver sensitivity	-80 dBm
Antenna Gain	0 dB
Noise temperature	290 K
System loss factor	1

5.5 Simulation Model

For our simulative studies, we use the network simulator ns-2 [1]. Several changes have been necessary to set up a model as described above. Mainly two enhancements have been implemented:

- an IEEE 802.11b multi-rate model, and

- a VoIP application for G711 and G729 including above mentioned QoS constraints.

The multi-rate enhancement decides for the data rate dependent on the distance between two nodes. We've chosen this design since we want to model optimal rate decisions, i.e., investigate the influence of transmission at lower rates not taking into account capacity wastes due to rate adaptation algorithms.

A VoIP application can be modeled simply by using bidirectional CBR streams. Since the perceived QoS for each VoIP participant is of interest, constraint metrics need to be observed on an end-to-end basis independently for each VoIP user. This has been realized by a separate VoIP application. In this work, the duration of a voice call is assumed to be large compared to the simulated time. This leads to the simplification that a voice call is modeled with infinite length.

Two investigations have been conducted: First, the number of VoIP terminals at one data rate is increased separately. VoIP calls are started with a distance of 100 simulated seconds until the QoS constraints are violated for one call. This investigation shows the upper limit

TKN-05-04   Page 17

at each specific data rate. Results are compared with Garg et al. [8], who considered the maximum number of calls for G711 and G729 calls previously.

In the second run, we start a single call a lower rate and all following calls at one specific higher rate. This results in the reduction invoked by one call at lower rate and thus shows the costs of this rate-degraded transmission.

All simulations in this work have been performed with 40 independent repetitions to allow a statement about the statistical significance of our results. This leads to 2000 simulation runs in total for each considered codec.

For all simulation results, 95 percent confidence intervals have been calculated. Since confidence intervals have been very small for all presented results, all mean values are significantly different. For a simpler graphical representation, the intervals are not drawn in the result figures of the next two chapters.

　　　　TKN-05-04　　　　Page 18

Chapter 6

Number of G711-coded VoIP calls

6.1 Capacity at each Data Rate

The maximum number of VoIP calls at each single data rate is displayed in Figure 6.1 for packetizations ranging from 10 to 50 ms. With increments of the packetization value, the number of possible calls increases, too, since the packetization is the reciprocal of the packet generation rate, i.e., the less packets are generated, the more calls can be served in the WLAN cell. The largest increments due to packetizations occur at a data rate of 11 Mbps, where the number of calls approximately doubles if packetization values are doubled from 10 to 20 ms and 20 to 40 ms.
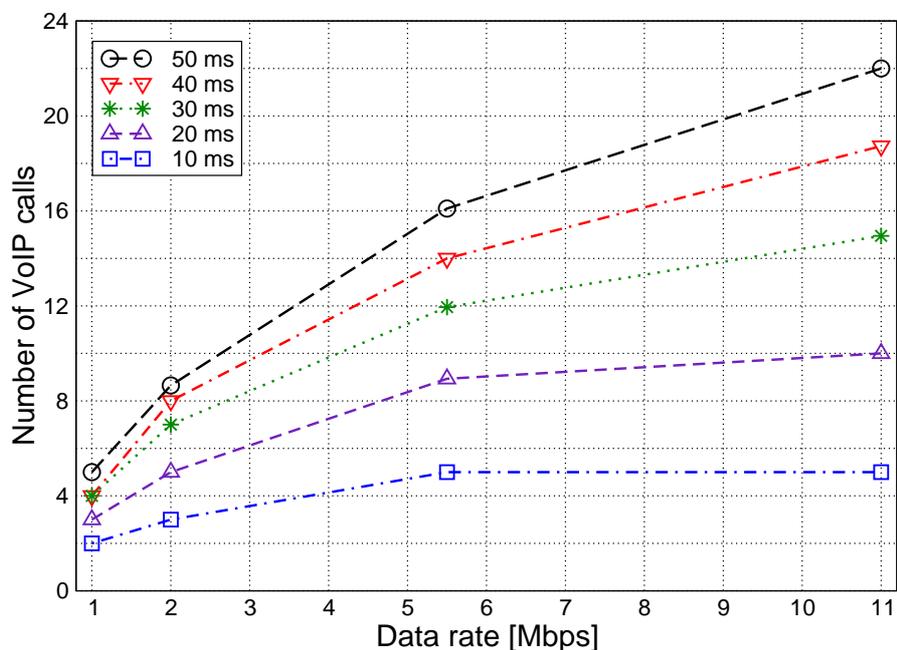
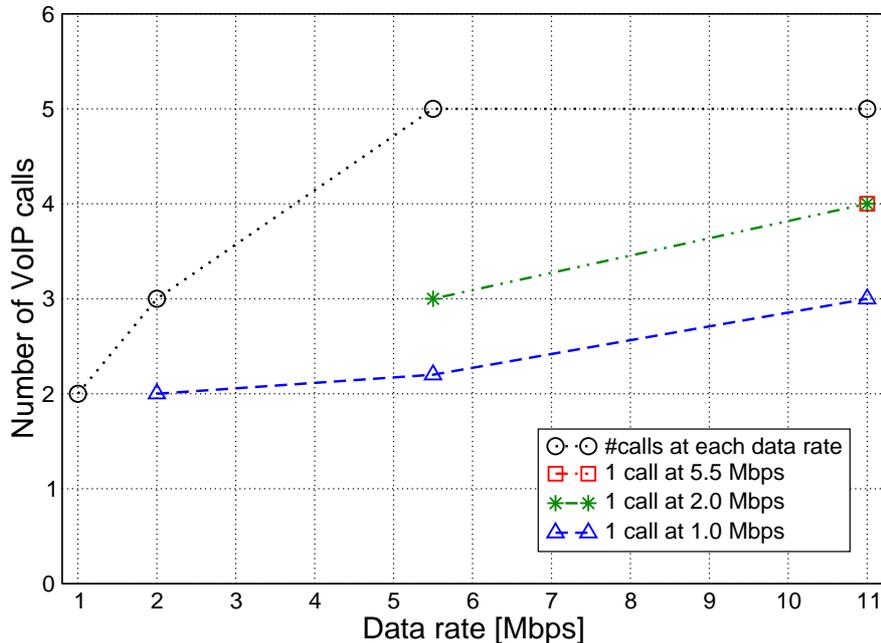Figure 6.1: G711, maximum number of calls at each single data rate

Figure 6.2: G711, 10 ms packetization, 80 Bytes voice payload

In combination with the data rates, the number of calls rises with increasing rates. For a packetization of 10 ms, no further increase is gained if the data rate is changed from 5.5 to 11 Mbps. At larger packetizations, the data rate has greater impact on the number of calls. This is a result of the packet size which increases for every 10 ms of packetization by 80 Bytes and—since it is part of the MPDU—is transmitted at the high data rate.

Garg et al. [8] determined the number of VoIP calls (for the same codec, various packetization values as well as without silence suppression) by considering numerically the occupancy time of each flow. At 11 Mbps, the authors gained 6, 12, 17, 21 and 25 calls for 10, 20, 30, 40, and 50 ms packetization. As shown in Figure 6.1, our simulations at a data rate of 11 Mbps enable 5, 10, 15, 19, and 22 calls. The small difference can be explained by the dynamic behavior that emerges when a new call is set up, e.g., the collision rate increases shortly until stations adapt to the new situation. This increases the delay of the VoIP packets temporarily. The whole WLAN cell, especially the downlink served by the AP, needs a certain amount of time to stabilize again. Since numerical flow investigations cannot take into account these dynamics, one to three more calls are gained by [8] dependent on the packetization.

6.2 Costs invoked by low-rate Transmissions

For packetization values ranging from 10 to 50 ms, Figures 6.2, 6.3, 6.4, 6.5, and 6.6 display the number of calls that are possible at higher rates in case that one lower-rate call is present already. The maximum number of calls at each single data rate is displayed additionally to allow a simple comparison with the results at each single data rate presented in the previous
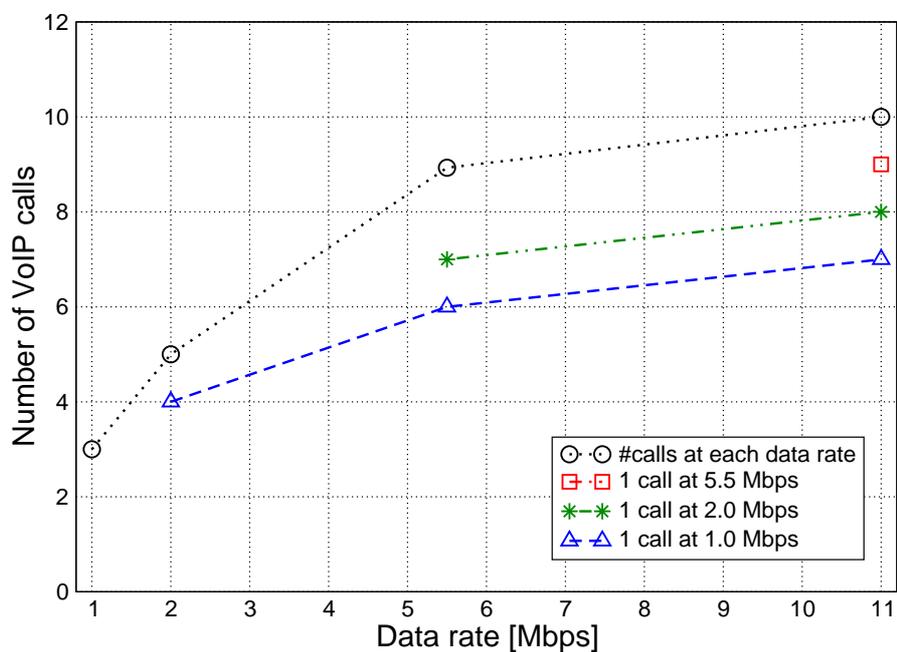
TKN-05-04   Page 20

Figure 6.3: G711, 20 ms packetization, 160 Bytes voice payload
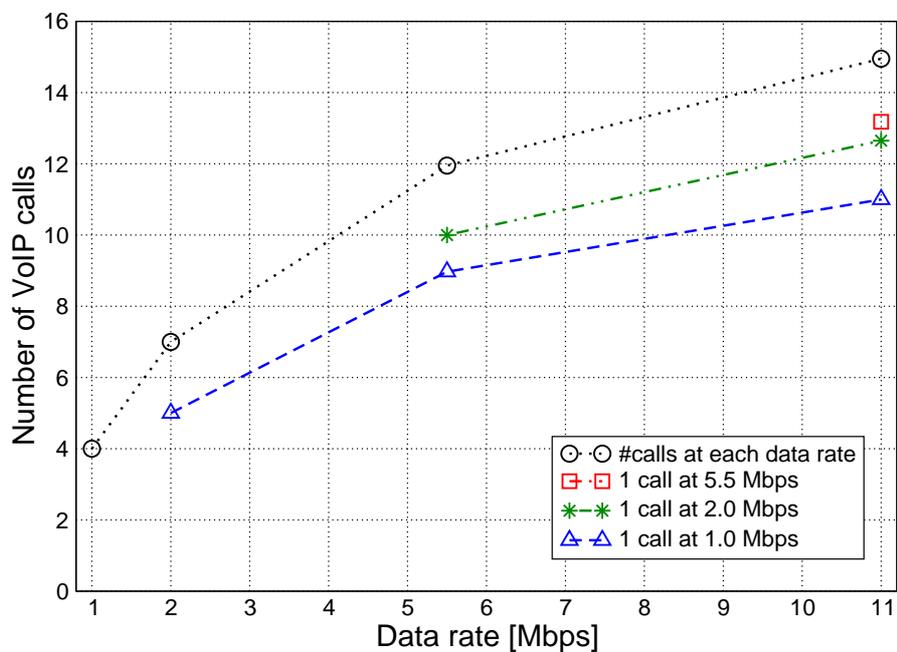

Figure 6.4: G711, 30 ms packetization, 240 Bytes voice payload
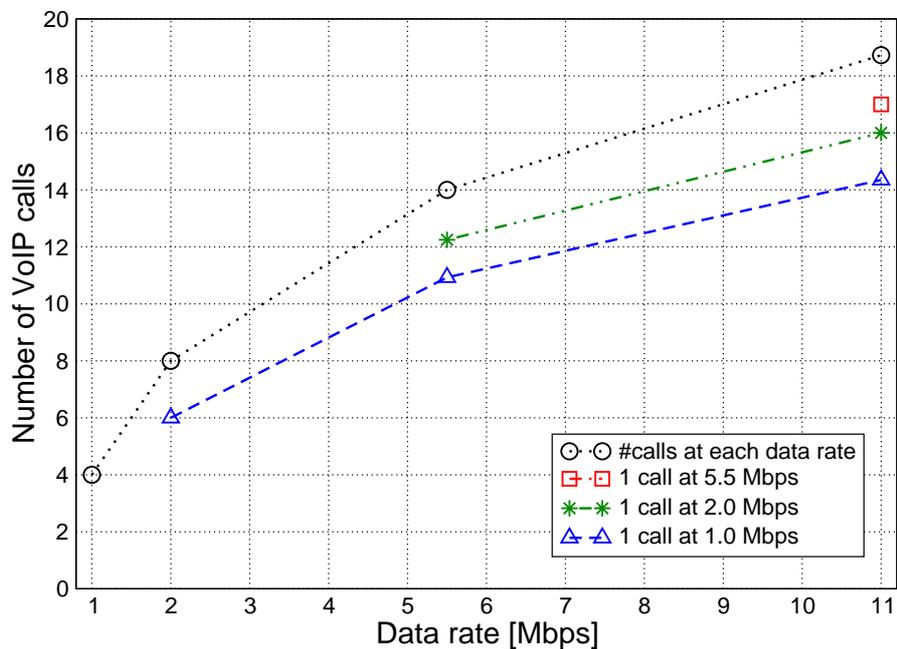
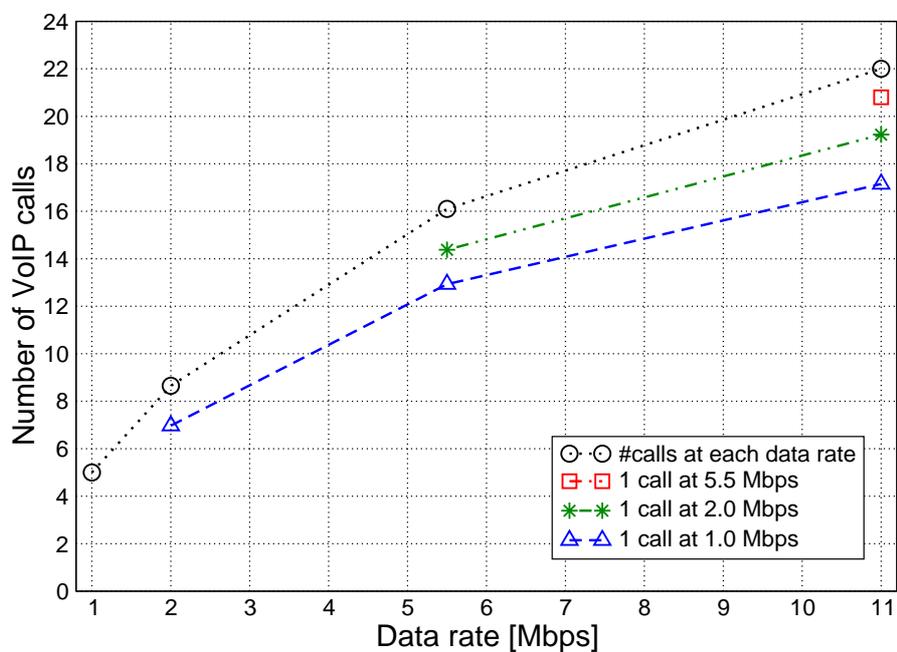Figure 6.5: G711, 40 ms packetization, 320 Bytes voice payload

Figure 6.6: G711, 50 ms packetization, 400 Bytes voice payload

section.

In the following, we explain the results shown in Figure 6.2 exemplarily for one low-rate call at 1 Mbps: With this single low-rate call, two other calls at a high rate of 2 Mbps are possible, thus no reduction in the total number of calls occur (compared to the results of Section 6.1, displayed by the "circled dots"). Applying a high rate of 5.5 Mbps instead, approximately two calls are possible plus the single user transmitting at 1 Mbps, i.e., the total number of calls is reduced by about two. With a data rate of 11 Mbps, three high-rate calls can be served. Together with the low-rate call, four calls in total are possible, which is a loss of one call.

In general, one call at 1 Mbps leads to the strongest degradation on calls at higher rates independent of the applied packetization value. Only at packetizations of 10 and 20 ms and a higher data rate of 2 Mbps the total number of calls is not influenced by the low-rate user. For all other packet-generation intervals and a high data rate of 2 Mbps, a 1-Mbps user decreases the number of calls by one. With larger rates of other users, reduction of calls appear also at packetizations of 10 and 20 ms.

Having one low-rate user at 2 Mbps results in no costs for 10 ms packetization and a high rate of 11 Mbps. All other combinations of packet generation intervals and higher data rates lead to reduced numbers of calls for a single 2 Mbps user.

For one user at 5.5 Mbps and all others at 11 Mbps, no costs occur with packetizations of 10 and 20 ms. At these values, the low-rate user seems to have no negative impact. This can be explained by a high generation rate of relatively small audio fragments of 80 and 160 Bytes, so that the impact of data rate itself is masked. For 30 and 40 ms, the increased packet sizes and the reduced packet generation rate leads to costs of one call, while again no costs appear for generation intervals of 50 ms. It seems that 20 packets per second with an audio payload of 400 Bytes do not have a large impact with respect to the 5.5 Mbps data rate of the single user, i.e., the impact of the low-rate transmissions is negligible due to the high number of calls at 11 Mbps.

Only absolute numbers of degraded calls have been considered so far. Since this does not give an idea about the impact of transmissions at lower rates, the difference between the number of maximum and lost calls is now set in relation to the number of maximum calls at each single rate. Table 6.1 lists the relations

$$\eta_{loss} = \frac{max_{single\,rate} - max_{mixed\,mode}}{max_{single\,rate}} \tag{6.1}$$

of lost to maximum number of calls for all rates and the five considered packetization values in percent. As expected, the transmissions at 1 Mbps lead to the strongest degradations of the number of calls. For packetizations of 10 and 20 ms, η_{loss} is higher at a data rate of high-modulation users equal to 5.5 Mbps than for 11 Mbps. All other values increase with the data rate of high-rate as well as low-rate user(s).

Table 6.1: G711, 10 to 50 ms packetizations, loss of calls in percent for one low-rate user

G711 packetization	data rate of high-rate users	data rate of single low-rate user		
		1.0 Mbps	2.0 Mbps	5.5 Mbps
	2.0 Mbps	0	–	–
10 ms	5.5 Mbps	36.0	20.0	–
	11.0 Mbps	20.0	0	0
	2.0 Mbps	0	–	–
20 ms	5.5 Mbps	21.6	10.4	–
	11.0 Mbps	20.0	10.0	0
	2.0 Mbps	14.3	–	–
30 ms	5.5 Mbps	16. 6	7.9	–
	11.0 Mbps	19 .7	8.7	5.2
	2.0 Mbps	12.5	–	–
40 ms	5.5 Mbps	14.8	5.4	–
	11.0 Mbps	18.0	9.2	3.9
	2.0 Mbps	7.9	–	–
50 ms	5.5 Mbps	13.5	4.5	–
	11.0 Mbps	17.5	8.0	0.9

Chapter 7

Number of G729-coded VoIP calls

7.1   Capacity of each Data Rate

For the G729 codec, Figure 7.1 shows the maximum number of VoIP calls that can be gained if all stations use one data rate only. With 11 Mbps, 6, 12, 18, approximately 24 and 29 calls are possible with a packetization of 10, 20, 30, 40, and 50 ms, respectively. As already discussed in Chapter 6.1, the number of calls increases with the packetization value and the data rate.
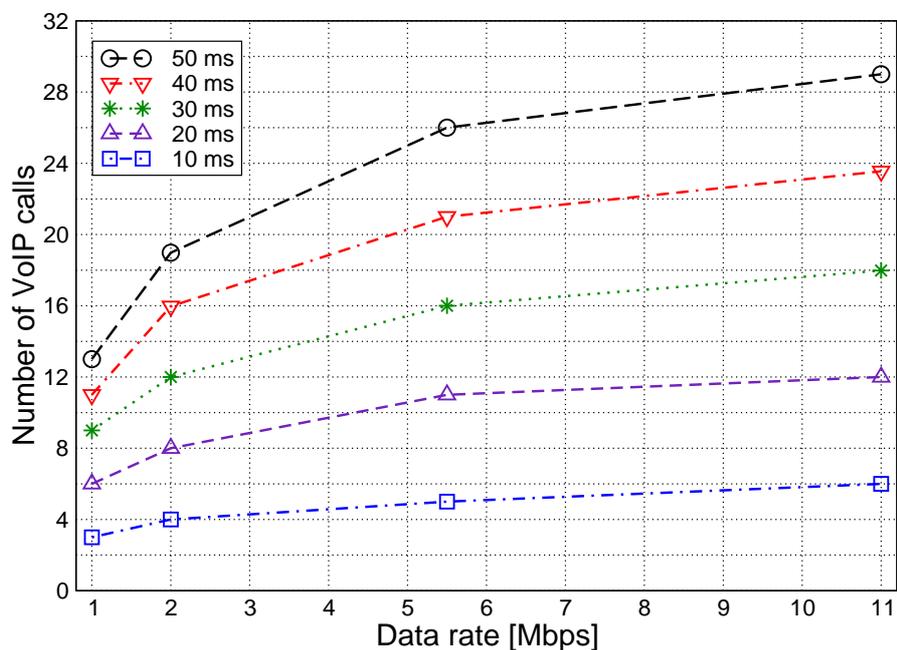

Figure 7.1: G729, maximum number of calls at each data rate

Again, we compare the maximum number of calls with the numerical results of Garg et al. [8]. For G729, the authors are able to serve 7, 14, 21, 28, and 34 calls in an IEEE 802.11b
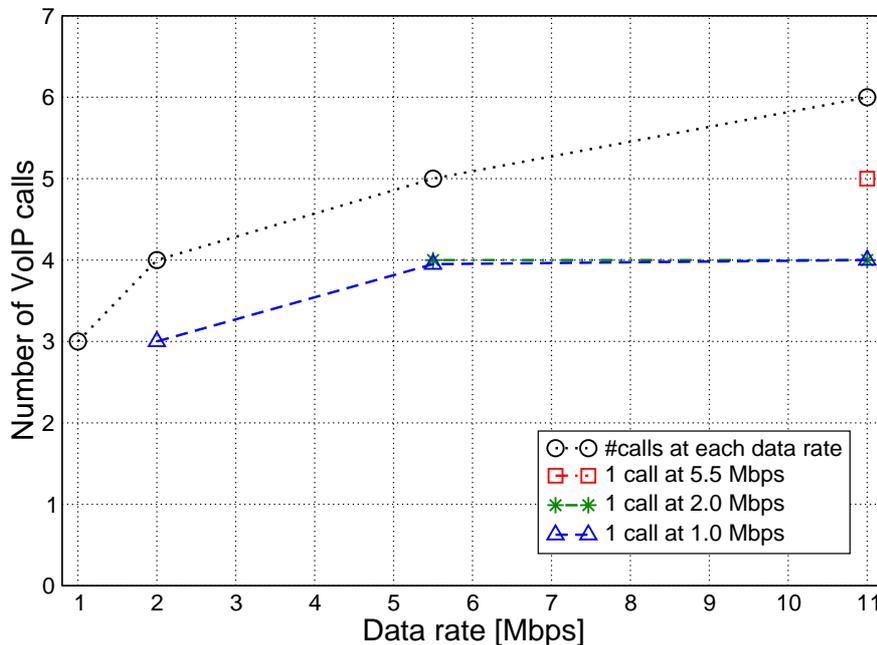
Figure 7.2: G729, 10 ms packetization, 10 Bytes voice payload

network for VoIP packets containing 10, 20, 30, 40, and 50 ms of audio data. Similar to the
G711 results, we explain the difference by the dynamic behavior that has not been considered
in the analytical investigation of [8]. However, the differences for G729 results are slightly
larger, since G729 enables in general more calls due to its compression which enlarges also
the gap between our and Garg's results.

7.2 Costs of low-rate Transmissions

Figures 7.2, 7.3, 7.4, 7.5, and 7.6 show the numbers of calls that are gained if a single user
has a VoIP call at a constant, lower data rate. Again, the maximum number of calls at each
single data rate is plotted to show the degrading impact of each low-rate call.

Having a low-rate call of 1.0 Mbps in combination with all other users at a higher rate of
2 Mbps leads only for a packetization of 40 ms to costs of one call. Changing the high data
rate to 5.5 Mbps results in no reduction for a generation interval of 10 ms, while one call less
is gained for all other packetization values. With 11 Mbps and packetizations of 10, 20, 30,
and 50 ms costs of one call occur, while the total number of calls is reduced by two in the
40 ms case.

Letting the single user transmit at 2.0 Mbps leads for 10 and 20 ms to the same behavior
as with the single 1 Mbps user. With 30 and 40 ms, no costs occur in combination with
5.5 Mbps users and a loss of one call with 11 Mbps. The low-rate user with 2 Mbps causes
no reduction of calls with a packetization of 50 ms.

With the single user at 5.5 Mbps, costs of one call only occur for a packetization value of
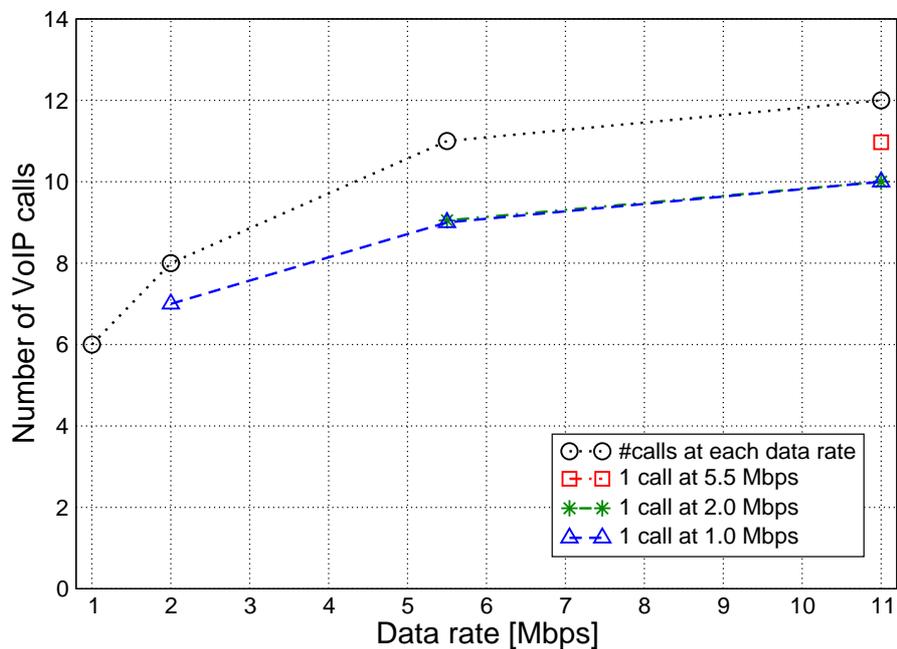
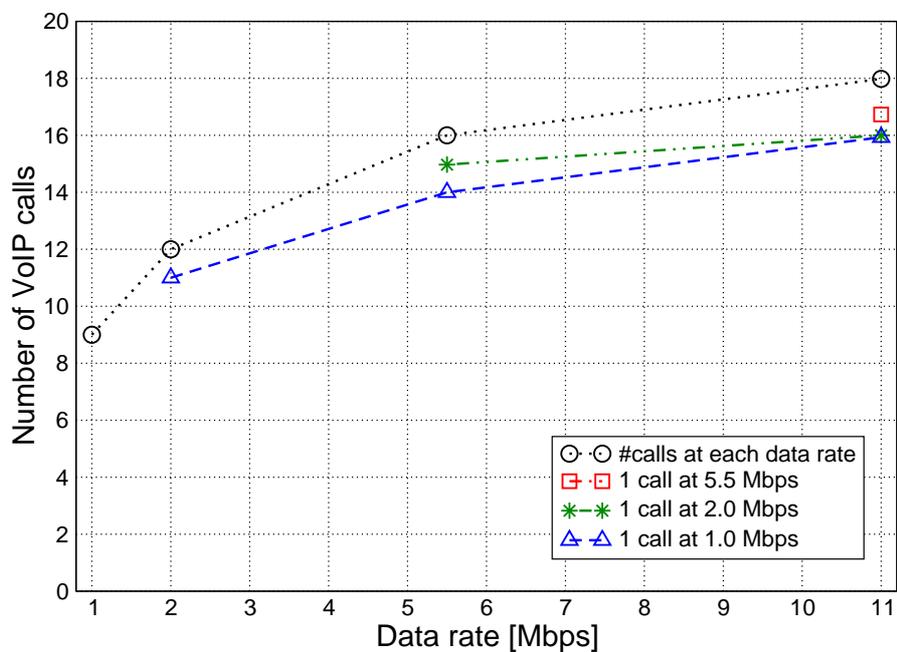Figure 7.3: G729, 20 ms packetization, 20 Bytes voice payload

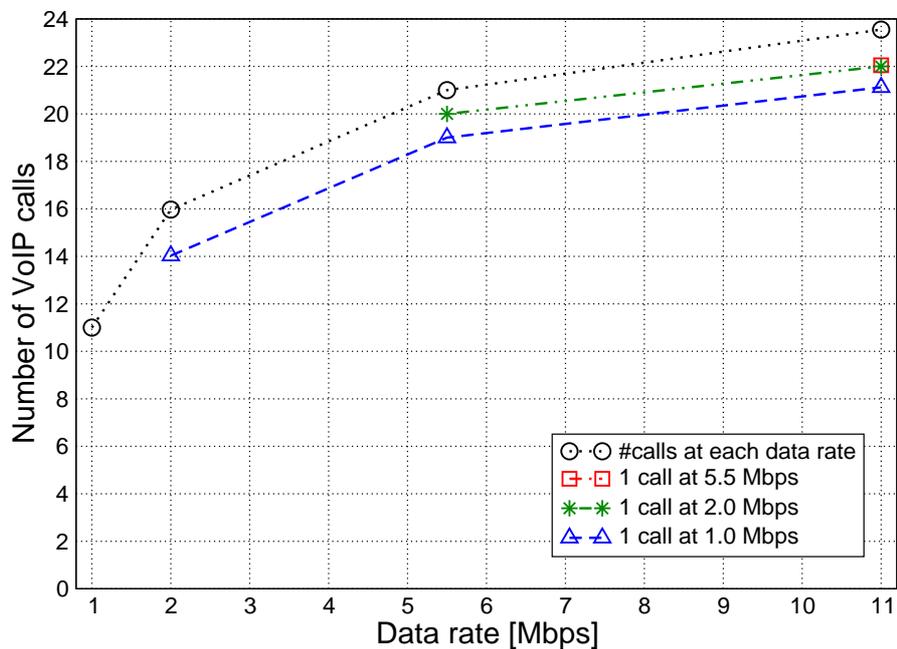Figure 7.4: G729, 30 ms packetization, 30 Bytes voice payload

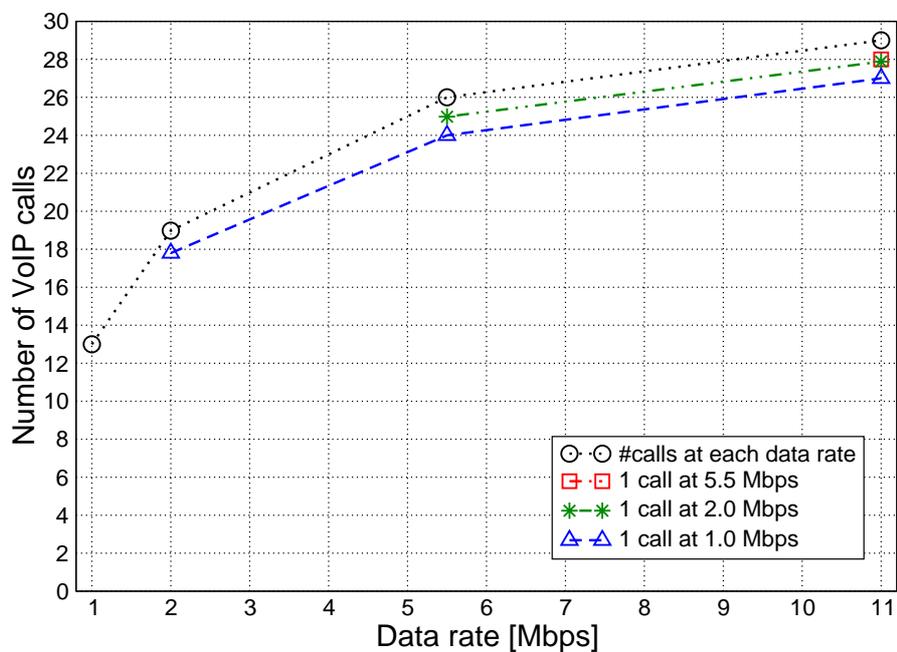Figure 7.5: G729, 40 ms packetization, 40 Bytes voice payload

Figure 7.6: G729, 50 ms packetization, 50 Bytes voice payload

TKN-05-04

Page 28

Table 7.1: G729, 10 to 50 ms packetizations, loss of calls in percent for one low-rate user

G729 packetization	data rate of high-rate users	data rate of single low-rate user		
		1.0 Mbps	2.0 Mbps	5.5 Mbps
10 ms	2.0 Mbps	0	–	–
	5.5 Mbps	0 (within CI)	0	–
	11.0 Mbps	16.7	16.7	0
20 ms	2.0 Mbps	0	–	–
	5.5 Mbps	9.1	8.6	–
	11.0 Mbps	8.3	8.3	0 (within CI)
30 ms	2.0 Mbps	0	–	–
	5.5 Mbps	6.3	0 (within CI)	–
	11.0 Mbps	5.8	5.5	1.4
40 ms	2.0 Mbps	5.9	–	–
	5.5 Mbps	4.8	0	–
	11.0 Mbps	6.1	2.3	2.1
50 ms	2.0 Mbps	0 (within CI)	–	–
	5.5 Mbps	3.8	0 (within CI)	–
	11.0 Mbps	3.4	0.4	0

40 ms and a higher data rate of 11 Mbps.

Similar to the G711 results, we list the relation of lost to maximum calls (Equation 6.1) for each data rate as well as for each considered packetization value for G729 in Table 7.1. If the relation is slightly larger than zero but the loss not significantly different, it is marked by the comment "within confidence interval (CI)".

The strongest losses occur for the smallest packetization with the low rates of 1 and 2 Mbps, which is a result of the relatively low capacity for single-rate transmissions. Overall, the reduction of calls rather decreases with increasing packetization values. This is due to the fast increasing number of calls at higher rates with increasing packetizations while only one single call is held at a lower rate. However, the further the packetization is increased, the more are losses affected by the lowest data rate of 1 Mbps which is responsible for the highest reduction at 20, 30, 40, and 50 ms. Finally, only the single call at 1 Mbps leads to a significant degradation at the highest generation interval of 50 ms, while all other low-rate calls are of minor impact.

Chapter 8

Virtual Utilization—Applications

8.1 Effect of Low-rate Transmissions

For every VoIP simulation of the two preceding chapters, we traced the virtual utilization that has been discussed in Section 4.1. A constant interval of 0.5 seconds has been used for the calculation of the average values. Figure 8.1 shows the utilization for a single 1 Mbps VoIP call and seven 11 Mbps calls over the simulated time.
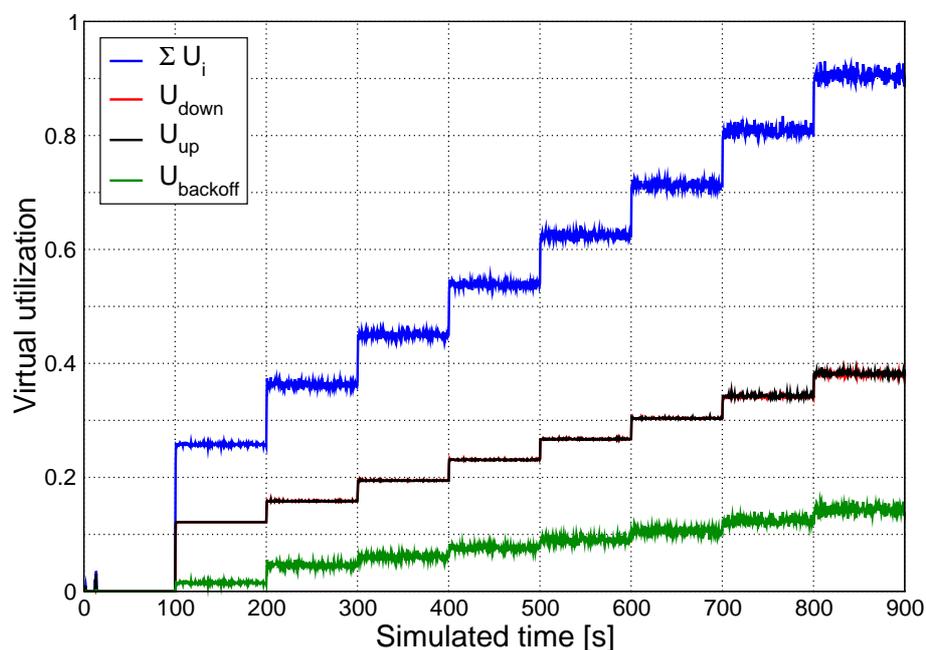

Figure 8.1: Virtual utilization at AP for G711-coded VoIP traffic, 20 ms packetization, 160 Bytes audio payload

The first call, which is initiated by the low-rate user, starts around 100 seconds. U_{up} as well as U_{down} are around 12 percent, while $U_{backoff}$ ranges from around 1.2 to 1.8 percent.

This results in a total virtual utilization U_{AP}, which is the sum of up-, downlink, and backoff utilization, of around 25 to 26 percent.

At 200 simulated seconds, the first high-rate user establishes its call which increases U_{up} respective U_{down} by 3 to 4 percent. In total, every additional high-rate call raises the virtual utilization by around 6 to 8 percent plus backoff utilization whose variance increases with the number of calls. After establishing the seventh high-rate call, the virtual utilization is already around 90 percent.

If another next high-rate call would be started, an overload case in the WLAN cell occurs. Due to to this, the VoIP QoS constraints of already established calls will be violated so that the call quality would be lousy.

8.2 Virtual Utilization for Admission and Handover Control

The virtual utilization is able to reflect the current situation in a WLAN cell and allows additional insights due to the separation in uplink, downlink, and overhead utilization. With this, we are able to identify users who consume a high amount of resources due to their current situation with respect to

- their traffic pattern,

- the channel characteristics, and

- the influence of the first two points on MAC and PHY protocols.

Combing the virtual utilization with specific load margins can be beneficial in certain fields. For example, Admission Control algorithms can be based on the virtual utilization criterion in combination with utilization margins. Considering the case shown in Section 8.1, the maximum value of the utilization would be around 95 percent. Above this limit, one would settle the "red" area, which should never be accessed since the overload will effect all ongoing VoIP sessions negatively. Reserving a certain amount of the resources for the dynamics occurring in a WLAN cell, e.g., due to channel characteristics and user mobility, leads to further gradation of the utilization margins. A challenging aspect is the approximation of these dynamics.

Being able to estimate the fraction of utilization that a new user evokes allows decisions whether to grant or to deny access to the network. Applying different margins dependent on traffic categories and perceived quality of the user allows fine-grained Admission Control decisions also in networks serving heterogeneous traffic.

Having the situation that a user is within the coverage area of several BSSs or has currently access to other radio access technologies leads to continuative aspects. With the virtual utilization, one is able to identify the user with the strongest negative impact, e.g., due to transmissions at lower data rates or a lot retransmissions. Thus not only the influence on the whole cell but also the impact on every single user (due to the combination with utilization margins) is assessable. As a result, the "best user to hand over" to another WLAN cell or radio access technology can be identified.

From the provider perspective, the virtual utilization enables the possibility to control the load in overlapping homogeneous as well as heterogeneous cells. Having users within

the coverage of multiple cells, Admission Control as well Handover decisions for the optimal distribution of the load are possible. Since the virtual utilization takes into account all overheads of the IEEE 802.11 MAC and PHY, it is an adequate measure for this.

Chapter 9

Conclusions

Firstly, this work discussed the overhead that arises out of IEEE 802.11 MAC and PHY and showed its large impact on VoIP traffic degrading the possible number of calls in a WLAN cell tremendously. To be able to evaluate the amount of resources that a certain traffic flow consumes, we motivate a measure—denoted as virtual utilization—that is based on the occupancy time for each transmission including inter-frame-spaces, (post)backoff, as well as durations of the data frame and the immediate ACK "on air". We placed the calculation of the virtual utilization at the AP, which has the reason that the AP is the bottleneck of a WLAN cell, thus its utilization state influences all users. With this, the backoff utilization of the AP (for transmissions in the downlink direction) has been included allowing a fine-grained evaluation of the situation in the WLAN cell.

In a next step, we considered a WLAN cell comprising of different modulation regions and studied the influence of transmissions at lower-rates in a pure VoIP network simulatively. Two codecs, G711 and G729 have been applied together with packetizations ranging from 10 to 50 ms. At first, the maximum possible number of calls for every codec have been identified for each single data rate of 1, 2, 5.5, and 11 Mbps. Since G729 compresses speech by the factor of eight (compared to G711), in general more calls are possible with this codec. These results have been discussed in the context of results published by Garg et al. [8] who determined the maximum number of calls via numerical investigations for an IEEE 802.11b network with a single data rate of 11 Mbps.

Afterwards, we determined the number of calls that is possible if one call is transmitted at a lower rate, while all other VoIP participants use a higher data rate. E.g., for G711 with a packetization of 20 ms, one call at 1 Mbps reduces the number of 11 Mbps calls by 3, i.e., the low-rate call has costs equal to two other high-rate calls. G729-coded speech with the same packetization and for the same example results in overall costs of one lost call. The lower call losses of G729 can be explained by the length of the audio payload, which is eight times smaller than G711 coded speech, and the higher number of high-rate calls in total thus decreasing the impact of the low-rate transmissions.

Finally, results of the virtual utilization have been shown for the case that one VoIP call is transmitted at 1 Mbps while all others use a data rate of 11 Mbps. The high consumption of resources for the low-rate call is directly identifiable by the virtual utilization. Thus the virtual utilization metric is an attractive candidate for Admission and Handover Control algorithms.

Acronyms

ACK Acknowledgment

AARF Adaptive ARF

ARF Automatic Rate Fallback

AP Access Point

BER Bit Error Rate

BSS Basic Service Set

CCK Complementary Code Keying

CFP Contention Free Period

CP Contention Period

CSMA/CA Carrier Sense Multiple Access with Collision Avoidance

CUE Channel Utilization Estimate

DCF Distributed Coordination Function

DIFS Distributed Interframe Space

HO Handover

IFS Inter-Frame Space

MPDU MAC Protocol Data Unit

MSDU MAC Service Data Unit

MT Mobile Terminal

PC Point Coordinator

PCF Point Coordination Function

PLC Packet Loss Concealment

PLCP Physical Layer Convergence Protocol

PPDU PLCP Protocol Data Unit

PSDU PLCP Service Data Unit

RBAR Receiver-Based Auto Rate

RSS Received Signal Strength

SIFS Short Interframe Space

ST Stationary Terminal

VHO Vertical Handover

RSSI Received Signal Strength Indicator

Bibliography

[1] The network simulator—ns-2. http://www.isi.edu/nsnam.

[2] IEEE Std. 802.11-1999. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, Reference number ISO/IEC 8802-11:1999(E). IEEE Std. 802.11, 1999 edition, 1999.

[3] IEEE Std. 802.11b 1999. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, Amendment 2: Higher-Speed Physical Layer (PHY) Extension in the 2.4 GHz Band, 2001.

[4] Jean-Pierre Ebert. *Energy-efficient Communication in Ad Hoc Wireless Local Area Networks*. PhD thesis, Techical University of Berlin, April 2004.

[5] M. Elaoud and P. Agrawal. Voice capacity in IEEE 802.11 networks. In *Personal, Indoor and Mobile Radio Communications, 2004. PIMRC 2004. 15th IEEE International Symposium on*, volume 1, pages 78– 82, 5-8 Sept 2004.

[6] S. Garg and M. Kappes. Admission control for VoIP traffic in IEEE 802.11 networks. In *Global Telecommunications Conference, 2003. GLOBECOM '03. IEEE*, volume 6, pages 3514–3518 vol.6, 2003.

[7] S. Garg and M. Kappes. An experimental study of throughput for UDP and VoIP traffic in IEEE 802.11b networks. In *Wireless Communications and Networking, 2003. WCNC 2003. 2003 IEEE*, volume 3, pages 1748–1753, March 2003.

[8] S. Garg and M. Kappes. Can I add a VoIP call? In *International Conference on Communications, 2003. ICC '03. IEEE*, volume 2, pages 779–783, 11-15 May 2003.

[9] C. Hoene, S. Wiethölter, and A. Wolisz. Predicting the perceptual service quality using a trace of voip packets. In *Proc. of Fifth International Workshop on Quality of future Internet Services (QofIS'04)*, Barcelona, Spain, September 2004.

[10] G. Holland, N. H. Vaidya, and P. Bahl. A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks. In *Seventh Annual ACM International Conference on Mobile Computing and Networking (MOBICOM)*, July 2001.

[11] Texas Instruments. *Voice over Packet White Paper*, January 1998. http://focus.ti.com/pdfs/bcg/vop_white_paper.pdf.

[12] Intel. *White Paper: Overcoming Barriers to High-Quality Voice over IP Deployments.* http://www.intel.com/network/csp/pdf/8539.htm.

[13] ITU-T. Recommendation G.107. The E-model, a Computational Model for Use in Transmission Planning, May 2000.

[14] A. Kamerman and L. Monteban. WaveLAN-II: A High-performance wireless LAN for the unlicensed band. *Bell Lab Technical Journal*, pages 118–133, Summer 1997.

[15] A. Koepsel, J.-P. Ebert, and A. Wolisz. A Performance Comparison of Point and Distributed Coordination Function of an IEEE 802.11 WLAN in the Presence of Real-Time Requirements. In *Proc. of 7th Intl. Workshop on Mobile Multimedia Communications (MoMuC2000)*, October 2000.

[16] A. Koepsel and A. Wolisz. Voice Transmission in an IEEE 802.11 WLAN Based Access Network. In *In Proc. of WoWMoM 2001*, pages 24–33, July 2001.

[17] Mathieu Lacage, Mohammad Hossein Manshaei, and Thierry Turletti. IEEE 802.11 rate adaptation: a practical approach. In *MSWiM '04: Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, pages 126–134, New York, NY, USA, 2004. ACM Press.

[18] K. Medepalli, P. Gopalakrishnan, D. Famolari, and T. Kodama. Voice capacity of IEEE 802.11b, 802.11a and 802.11g wireless LANs. In *Global Telecommunications Conference, 2004. GLOBECOM '04. IEEE*, volume 3, pages 1549–1553, Nov 2004.

[19] K. Medepalli, P. Gopalakrishnan, D. Famolari, and T. Kodama. Voice capacity of IEEE 802.11b and 802.11a wireless LANs in the presence of channel errors and different user data rates. In *Vehicular Technology Conference, 2004. VTC2004-Fall. 2004 IEEE 60th*, volume 6, pages 4543–4547, September 2004.

[20] Jd. P. Pavon and Sunghyun Choi. Link adaptation strategy for IEEE 802.11 WLAN via received signal strength measurement. In *IEEE International Conference on Communications, 2003. ICC '03.*, volume 2, pages 1108–1113, May 2003.

[21] John G. Proakis. *Digital Communications.* Irwin/McGraw-Hill, 2001.

[22] Theodor S. Rappaport. *Wireless Communications—Principles and Practice.* Prentice Hall, 2002.