Rapport de stage d'option scientifique

Multicast Video Transmission Over Wireless LANs

Wolf H. Lauppe
Ecole Polytechnique
Promotion 2002

NON CONFIDENTIEL

Option : Département d'informatique
Champ de l'option : Réseaux et Télécommunications
Directeur d'option : Gilles Dowek
Directeur de stage : Thierry Turletti
Dates du stage : 11 avril 2005 - 31 juillet 2005
Adresse de l'organisme :
Projèt Planete
INRIA Sophia Antipolis
2004 route des Lucioles
06902 Sophia-Antipolis

Abstract

multicast video transmissions in wireless LANs encounters a number of difficulties not solved today. This report gives an overview of the challenges multicast video transmissions have to tackle in wireless LANs. We stress on a realistic wireless LAN packet erasure model, a major requisite to evaluate multicast transmission solutions. A new discrete event network simulator framework which incorporates this model is presented. We outline a solution of jointly optimized communication layers and devise and evaluate an enhanced MAC layer for use in such an architecture.

Résumé

La transmission vidéo en multipoint sur réseaux sans fil IEEE 802.11 pose aujour'hui d'intéressant problèmes de recherche. Ce rapport de stage donne une vue d'ensemble des problèmes à traiter. Il est indispensable d'utiliser un modèle de canal réaliste pour les réseaux IEEE 802.11. On propose un nouveau simulateur a événements discrèts qui incorpore le modèle proposé. On expose notre solution d'optimisation conjointe de différentes couches de communication et on évalue une couche MAC améliorée pour ce type d'architecture.

Contents

Chapter 1

Introduction

802.11 wireless local area networks are one of the fastest growing network technologies in the wireless communications field. Large scale effects leading to quickly dropping prices and establishing wireless LAN as an inexpensive method of wireless connection have contributed to an incredible success in the last years. Computer and communications industries have embraced this technology. For mobile devices such as Personal Digital Assistants (PDAs) and Laptops it has become a default wireless technology equipment.
With wireless networks gaining prominence and acceptance, it is foreseeable that streaming of video will be a critical part of the wireless LAN infrastructure.
More and more public places are covered by hotspots. This offers new possibilities: Travelers at an airport, for example, could use a laptop or PDA to log into a wireless video service and watch a television broadcast.
The use of wireless LANs as a means of broadcasting video may be fueled by the fact that wireless LAN technology is about to be deployed in public transports: Today air planes start being equipped with wireless LAN technology. Wireless LAN deployment for trains is now in the test stage.
The examples show: wireless LAN establishes itself as the dominant transmission technology at all places people are waiting and most likely use video broadcast as means of information and entertainment.

In all these scenarios, *multicasting* video to a group of receivers might be the answer to achieve high throughput and video quality. Multicasting video, instead of sending each receiver a separate stream, results in a much more efficient use of the shared media.

1.1 Challenges

While the examples above illustrate the need for a good support of multicast video in wireless LANs, wireless multicast transmission encounters a number of difficulties, which need to be solved, before it can be used effectively.

The current standard defines that packets transmitted via multicast are not acknowledged (see 2.3). Several drawbacks are associated:

- The contention window is not adaptive to congestion. Multicast sends blindly packets. This is a risk for the stability of the protocol in high load conditions and not a fair solution regarding other traffic.
- There is no error correction at the MAC layer. In unicast transmissions errors induced by the wireless medium are reduced by the usage of RTS/CTS protection, packet fragmentation and retransmission mechanisms. In multicast transmission the error rate is directly exposed to upper layers. This means that standard video streaming solutions fail in achieving a satisfactory result of video quality.
- The physical transmission rate cannot be adapted to the channel conditions. There is no feedback for the sender to chose the rate. Packets have to be sent at the basic rate which reduces the throughput of multicast considerably.

In wireless LANs multicast video transmissions struggle also with the great heterogeneity among receivers. Mobile devices cover a wide range of sizes from laptop to small PDAs. They might have different screen sizes, transmission power limitations, and different processing capabilities. Besides receivers can experience distinct channel conditions depending on their location with differing error rates which adds to the heterogeneity. This makes a one size fits all multicast video diffusion impossible. One approach to handle this heterogeneity is to use layered or hierarchical encoding.

Wireless LAN channels are not only heterogeneous between the receivers in one area. They also exhibit a high variability with time. In contrast to wired networks, channel conditions of wireless networks exhibit high fluctuations over time and can suddenly change. The fast changes make the task difficult to group the receiver in different layers. A solution adaptive to the quick changes of the media is needed.

In this report we consider methods to improve the current situation of wireless LAN video transmissions. To overcome the hurdles, an approach including communication between the different layers of the layered network architecture is suggested. We propose and evaluate an adapted MAC layer including a physical rate selection algorithm for multicast. A layered multicast video transmission scheme is outlined as a solution. In order to evaluate the performance of the new layer a network simulator with a sophisticated loss model that describes adequately the packet losses caused by the physical conditions of a wireless of a wireless LAN network is needed. We develop a

discrete event simulator, able to show the performance of the new mac layer under various network loads.
This report is organized as follows: In the next chapter we give a overview of the state of the art. Chapter 3 explains the envisioned MAC layer encompassing the rate selection algorithm in detail. Chapter 4 explains the architecture of the new network simulator. In Chapter 5, the new MAC layer is evaluated. Chapter 6 finally draws conclusions and gives an outlook on further work.

Chapter 2

Background and related work

2.1 Real-time Transport Protocol (RTP)

The key standard for audio and video transport in IP networks is the Real-time Transport Protocol (RTP) [1, 2] along with its associated profiles and payload formats. RTP aims to provide services useful for the transport of real-time media. The services include timing, recovery loss detection and correction payload and source identification, reception quality feedback, media synchronization and membership management. The associated RTCP protocol, which provides reception quality reporting of the receivers together with RTP makes available the services needed for an layered video approach to work. RTP is mostly used over UDP. This enables the application layer to have complete congestion and error control, adapted to the realtime requirements of the data transmitted. Congestion and error control can also be specially adapted to the high error rate specific to wireless LANs.

2.2 Layered coding

As described transmission of multimedia flows over multicast channels is confronted with the receivers heterogeneity problem. Layered hierarchical coding is often proposed as a solution to overcome this problem [3–5].
In this approach, the video is encoded in a base layer and one or more enhancement layers. The base layer can be independently decoded, but the enhancement layers can be decoded cumulatively (i.e. layer k can only be be decoded along with layer 1 to the layer k-1). The enhancement layers contribute to the improvement of the video quality and lead to a progressive refinement of the signal. Based on their needs and their wireless channel reception capabilities stations receive additionally to the base layer one or more enhancement layers.

An algorithm is needed to classify the receivers into multicast groups and to manage and choose the optimal the number and bandwidth of the layers. The Source-channel Adaptive Rate Control (SARC) [5] can be used to perform this task.

SARC adapts dynamically the number of layers their bandwidth and FEC redundancy according to RTCP feedback sent periodically.
At the beginning of each round of adaptation the source announces the number of layers and their respective rates via RTCP sender reports. Each source layer is transmitted to an IP multicast group.
Each receiver measures network parameters. Estimated bandwidth and loss rates are then conveyed to the sender via RTCP receiver reports.
Receivers are classified according to similar reception behaviours.
The source then proceeds with a dynamic adaptation of the number of layers and of their rates in order to maximize the quality perceived by the different clusters.

For the receivers to receive video with high quality in the presence of the high error rates of the wireless medium it is necessary to exert some form of error control.

While choosing the bandwidth of the layers the algorithm must also choose dynamically the error correction redundancy to be added the different layers. Various schemes have been proposed for error correction including FEC (forward error correction), ARQ (automatic repeat request) and hybrid approaches [4] [5].

2.3 IEEE 802.11 Standard

The IEEE 802.11 standards [6,7] use the same logical link layer as other 802-series (including the popular 802.3 wired Ethernet standard) and use compatible 48-bit hardware Ethernet addresses to simplify routing between wired and wireless networks. Unicast and multicast is supported. The Medium Access Control and Physical layers are newly defined reflecting the different properties of wireless networks:

The IEEE 802.11 standards specify the Medium Access Control (MAC) layer, as well as different physical (PHY) layers. Currently for the MAC layer the standard defines two medium access coordination functions: the contention-based Distributed Coordination Function (DCF) and the contention free based Point Coordination Function (PCF). We only consider the DCF access method. The PCF access method is not mandatory and therefore is rarely implemented in current products. The DCF access method is based on the Carrier Sense multiple Access with Collision Avoidance (CSMA/CA) principle. Each STA has the same priority when competing for an empty slot time. Before an STA attempts a first packet transmission it has to sense the medium. If the medium is found idle equal to the Distributed Inter Frame Space (DIFS) the packet will be transmitted directly. Otherwise the STA enters into backoff and randomly sets its backoff timer within the range of the Contention Window. The backoff timer is decremented by one very slot time the medium is sensed to be idle and it is frozen when the medium is sensed busy. If the counter reaches zero the packet is transmitted.

2.3.1 Unicast transmissions

Upon the correct receipt of a packet, the receiver has to send an acknowledgment (ACK) after a time equal to the Short Inter Frame Space (SIFS). If no ACK is received the sending STA assumes a collision, doubles its current CW, chooses randomly a value between [0,CW] and retransmits the packet when the timer reaches again zero.
For long packets, exceeding the RTS-threshold the standard defines a RTS/CTS procedure to be used. A short Request To Send (RTS) frame is sent, which is after the time of a SIFS acquitted by the receiver by a Confirm To Send frame (CTS). After the successful receipt the sender waits for the time of a SIFS and transmits the data packet. The correct transmission is confirmed by an ACK. Together with the virtual network reservation functions this reduces the probability of collisions (figure 2.1).

2.3.2 Multicast transmissions

In multicast the sender addresses a group of receivers. The standard proposes a simple solution to the question who of the receivers (> 1) should give the feedback: There is no feedback. Like in unicast after the medium is found idle equal to the Distributed Inter Frame Space (DIFS) the packet will be sent directly, otherwise after setting and counting down the backoff timer. Because there is no ACK the sender sends blindly and cannot adjust the Contention Window size after a collision. The MAC layer cannot decide whether a packet needs to be retransmitted and the high packet error rate, even higher due to the absence of the RTS/CTS is directly exposed to the upper layers. Absence of feedback frames (CTS, ACK) has another disadvantage: The physical rate 2.4.1 cannot be adapted to the physical channel conditions.

2.4 Packet erasure channel models for wireless LANs

Wireless LAN channels exhibit peculiar channel characteristics. The distinct nature of the wireless medium induces packet loss patterns that are very different from other IP networks: Due to reflection, refraction and scattering of radio waves by surrounding terrain the transmitted signal most often reaches the receiver by more than one path, resulting in a phenomenon known as *multipath fading*. The signal components arriving from indirect paths and the direct path (if it exists) combine and produce a distorted version of the transmitted signal. In narrow-band transmission the multipath medium

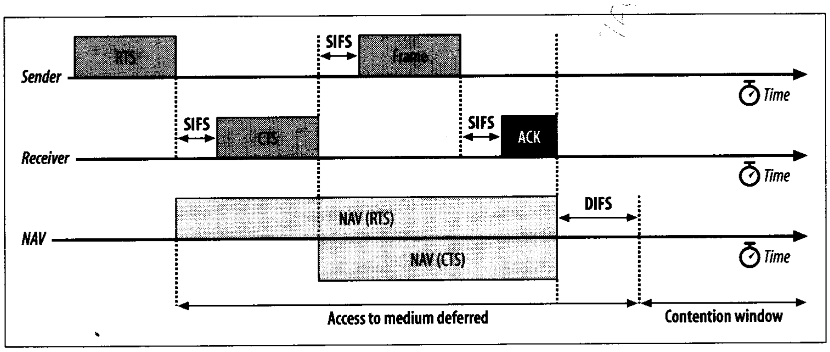

Figure 3-5. Using the NAV for virtual carrier sensing

Figure 2.1: DCF virtual carrier sensing, for packets transmitted in unicast

causes fluctuations in the received signal envelope and phase. In wide-band pulse transmissions on the other hand, the effect is to produce a series of delayed and attenuated pulses (echoes) for each transmitted pulse. The constructive and destructive combination of these pulses excite fluctuations in the received signal.

Often packet erasures are assumed to be independent and identically distributed *(i.i.d)*. This model is used by many network simulators. Though this might be a reasonable approximation valid for most IP networks it does not adequately reflect the statistical properties of packet loss in a wireless LAN.
A more sophisticated mathematical model that reflects the most important characteristics of the transmission medium is a Poisson process P_{t_i} with Rayleigh distribution [8, 9].
The channel is modeled as a process with discrete sample space. The value P_{t_i} is the Signal to noise ratio (SNR) of the channel for the time interval $(t_i, t_{i+1}]$. The time instants $t_1 < t_2 < t_3 < ... < t_i$ have a mean distance T_c which is known as the stationary time.
Whether a packet can be received correctly depends also on the transmission rate the packet is sent: Every transmission rate defines a transmission mode for the symbols and may have different internal FEC protection [7]. Therefore a specific target Signal to Noise ratio SNR^* is required for the transmission to succeed. If during the transmission of a packet the SNR of the channel drops below SNR^* the model assumes, that the received packet at the receiver side has to many bit errors to be corrected by the internal FEC. The CRC-check fails and the packet is dropped.

It is important to note that due to the layered nature of network protocols, corrupted packets are not available for the application layer, even if some parts are correct and the packet has some value of information. The wireless channel appears as a *packet erasure channel*. Error correction methods on the application layer have to take this into account.

2.4.1 Rate selection

The IEEE 802.11b specifications for the PHY layer define different transmission modes. IEEE 802.11b defines a 2.4-GHz spread-spectrum wireless LAN capable of operation at bit rates of 1,2,5,5 and 11 Mbit/s using BPSK, QPSK and CCK modulation.

The standard does not define how to choose the best rate for transmission. Control frames are exchanged with the basic transmission rate. Choosing the physical transmission mode for data packets is left as an implementation detail.

Although often neglected the selection of the correct rate is critical point in reaching high throughput. With the big differences in transmission speed (between 1 and 11 Mbit/s) a wrong rate can degrade performance considerably.

Various rate selection algorithms have been proposed for wireless LANs. [10–12]. The difficulty for a rate selection algorithm is to achieve the trade-off between high resilience against transmission errors which is guaranteed by the lower transmission modes and the gain in throughput obtained by a transmitting at higher rate.
An algorithm which has shown good performance in unicast transmissions is the CLARA algorithm [12]. We give a short overview.
Clara is a dynamic rate selection algorithm. The sender gets feedback information from the receiving station. This information is transmitted through the use of reserved bits in the Service field of the Physical Layer Convergence Protocol (PLCP) of the frames send back by the receiver. Thus for the feedback no modification in the MAC frame format is necessary. By the use of these bits the feedback of the receiving station is piggy-backed through both CTS and ACK frames.
The advantages of this approach are that CLARA can better differentiate between packet loss due to a bad physical channel and packet loss due to collision. Other rate selection algorithms which choose the rate based on the history on the past received ACKs, are much slower at adapting the rate to the quickly changing channel conditions. If a transmission rate is chosen only on the information whether received or not it takes several frame exchanges until the best rate is found. If packets are lost due to collisions, such an algorithm yields suboptimal results because the transmission rate is incorrectly assumed to be to high and the rate lowered. The information in the PLCP bits of the Control Packets sent back by the Receiver (CTS, ACK) enable CLARA to quickly choose a good transmission mode for the momentary physical channel conditions.

If fragmentation is used, which is possible for unicast transmissions CLARA proposes to deduce the channel coherence time based on the feedback received, and to improve the performance further by adapting the length of the fragments to this parameter. A good selection of the fragment size results in an elevated probability for the channel to be stable for the complete transmission of the fragment. The transmitted fragments are less likely to be received corrupted. It is shown that the overhead of fragmentation can than outweighed by this effect and additional performance gains can be achieved.

2.5 Related work

As of today the paper [4] is the only paper available, explicitly treating multicast video transmission in wireless LANs. In this paper the authors consider changes at the application layer to improve performance. Layered coding with an ARQ/FEC error correction scheme is suggested. The performance of this error correction scheme is evaluated experimentally for one single receiver and proves to be a good approach. Optimizing video transmission for a multiple user case is only addressed theoretically. For the calculations, a model of uniform distributed packet errors is assumed. Because loss characteristics significantly differ in wireless LAN channels from the (i.i.d) model it questionable whether the results obtained are realistic. The paper does not propose a solution to adapt video transmission to the changing reception capabilities of the receivers.

Chapter 3

Our approach

3.1 Scenario

A setting which exemplifies the application of wireless LAN video broadcasting is transmitting video at a public hotspot in an airport: A group of wireless LAN equipped devices wish to receive a video stream. The receivers might be different devices (Laptop, PDA). They have various screen sizes and reception capabilities. (depending on the distance to the base station and antenna characteristics).

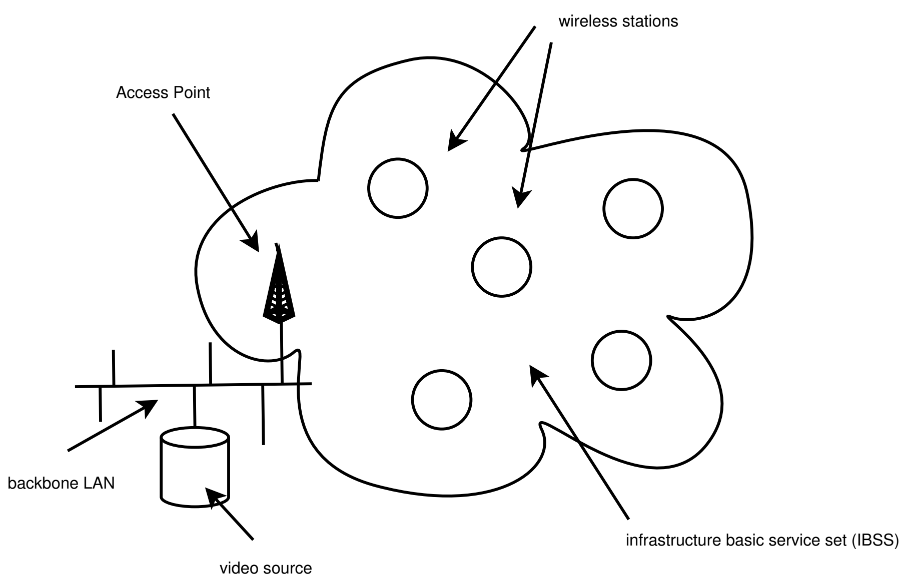

Figure 3.1: an access point multicasting video to different heterogeneous receivers

All receivers are assumed to be in the same basic service set. They receive the video multicasted by the access point and communicate in infrastructure mode. The video source is in the backbone. The solution envisioned aims to improve the transmission of the last hop, the wireless LAN which is likely to be the bottleneck of the whole video streaming system. The simulations therefore measure the performance of the wireless transmission, not taking into account delay and packet loss caused by the backbone.

The video is hierarchical encoded in multiple layers. A variant of the SARC algorithm manages the optimal allocation of bandwidth to the layers and the best FEC protection. The algorithm also assigns the receivers to different groups, each multicast group receiving one layer. Clustering of the receivers is done in feedback rounds. The RTP standard defines that the overhead of control packets should be less than 5%. In a group of about 30 receivers we can assume that a feedback round will take place every five seconds.

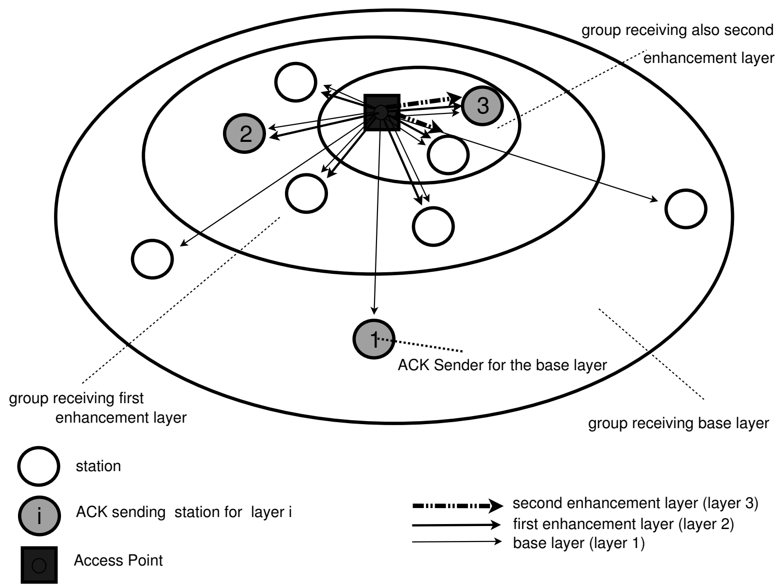

Figure 3.2: layered architecture

3.2 MAC layer enhancements

The wireless LAN 802.11 standard states that multicast packet transmissions are not acknowledged. We propose the following change:
In each video layer multicast group, one station is responsible to acknowledge each packet of the layer with an ACK-frame.
If the sender does not receive the ACK, the packet is retransmitted.

The ACK-frame also allows the responsible receiver to transmit feedback about the channel state. This establishes a basis for a rate selection algorithm to work. Using the CLARA-Algorithm based on the feedback of the ACK-sender the MAC chooses the physical rate for the next packet of this

layer to be transmitted.
In each feedback round the SARC algorithm chooses one station in the multicast group to be the dedicated ACKsender for this group. The information of the chosen ACKsender will be communicated to the MAC layer by the use of reserved bits in the packet headers.

Chapter 4

The *wirelessSim* network simulator

In this chapter we introduce *WirelessSim* the new network simulator we designed to measure the performance of the proposed MAC layer enhancements.

There exist a number of high quality simulation tools in widespread use. They have grown over the years and have become big toolboxes, supporting simulations with a sophisticated topology, and different level of detail with up to thousands of nodes. The venerable and widely used ns-2 simulator of the Vint Project [13] is one example of such a tool allowing to simulate networks in different abstractions. Regarding aspects such as the network topology they are very mature tools but they still lack good wireless LAN packet erasure models.

As it was clear that to backstitch the proposed channel packet erasure model into ns-2 would be difficult at best, and the strengths of ns-2 like a complex network topology are of no advantage in our simulations, we have decided to develop a new simple efficient simulator better adapted to our requirements.

WirelessSim has been implemented in object-oriented Java. Object-oriented design with clear abstractions allows the simulator to be quickly adapted to specific needs. It has proven to be fast, even while executing simulations using several gigabytes of precomputed channel data. Different simulations can be quickly constructed by combining basic building blocks. Several recording facilities allow fine grained control over logging of simulation events.

4.1 Discrete event simulators

Network simulating is a special kind of a physical system network simulation. In a physical system network, physical processes operate autonomously except to interact with other physical processes by the exchange of messages or events. The message sent by a (physical) process depends on the characteristics of the process (initial state, rules of operation).
A variable clock holds the time up to which the physical system as been simulated.
A data structure called the *event list* maintains a set of messages, with the associated times execution is scheduled for the future. At each step the message with the smallest associated future time is removed form the event list. Sending this message may in turn cause other messages to be sent in the future (which then are added to the event list). The clock is advanced to the time of the message transmission that was just calculated.
This form of simulation is called event-driven, because events (i.e. message transmission in the physical system are simulated chronologically and the simulation clock is advanced after simulation of an event to the time of the next event. There is another important simulation scheme; time-driven simulation in which the clock advances by one tick in every step and all events scheduled at that time are simulated. Because in its domain it is faster and yields the same accuracy wirlessSim uses the event-driven approach.

4.2 Design of the network simulator

The design closely matches the design of real network protocol stacks.

We present a brief overview of the architecture of the network simulator in this section.
In the design phase of a simulator it has to be decided up to which level of detail the real world is modeled. Our simulation is done at the MAC level of wireless LAN.

Synchronization is done by the simulation engine. It administrates the event list, executes the different events in the right order and advances the central clock time. It is at the same time the process which simulates the wireless medium. All events are actions on the MAC layer like waiting for a certain time or sending a frame. Events are evaluated for their implications (collisions, good transmissions) with the help of a channel package which provides the physical channel packet erasure model presented in section 2.4. This might entrain new events to be triggered in the future (setting NAV, packet dropped). These are created by the simulation engine.
There is no exchange of event messages between the other processes of the MAC layer, they all communicate via the simulation engine.
This design prevents deadlocks because it avoids circular event dependencies.

Other parts of the Simulator are the *MAC protocols*, *actions* (basic events on the MAC layer) and *traffic streams* (modeling together with some extra classes the upper layers). Finally a recorder facility collects the statistics.
A short overview over the different parts follows.

4.2.1 Simulation engine

The heart of the simulator is the *simulation engine*. It administrates the event list and executes the different events in the right order. For this purpose every MAC protocol has to be registered at the simulation engine. In every step the simulation engine determines the next action to be finished. The simulation engine sets the global clock time to the end time of this action, and if the action was created by a MAC protocol, sends the protocol event which summarizes what has happened and requests a new action.
The simulation engine is also responsible for updating the state of the network, seen by every receiver. If a transmission starts in a wireless LAN other stations do not see this immediately. The signal arrives after some time known as the propagation delay, depending on the distance. This means that every station has a local view of the state of the medium. The information whether the medium is busy or idle and the virtual reservations saved in the Network Allocation Vector have to be updated after this propagation time for every station who could perceive the transmission. To accomplish this at the beginning and the end of the action representing the sending of a packet, the simulation engine automatically creates actions needed together with their execution times. They are added to the event list and thus automatically executed at the right place in the stream of the pending events. This approach simplifies the design of the simulator. Keeping state of the medium seen by the receivers up to date uses the same uniform approach of actions as the other MAC releated events.
An object oriented approach allows to model the actions as objects which can implement its own logic. if the local network state changes they are asked by the simulation engine to recalculate the time when they are finished.
This reduces the complexity in the simulation angine and in the MAC protocols. Constructing MAC protocols is facilitated. Design, maintenance and verification of the MAC protocols is simplified.

4.2.2 MAC protocols

In section 2.3 we gave an overview, of the different protocols sending a packet (with and without RTS/CTS protection, fragmentation, and without ACK acknowledgment).

The protocols are modeled in the simulator as MAC protocols. All implement the MAC protocol interface which is a common interface and defines

the interaction to be used between the simulation engine and the protocol. The simulation engine and the MAC protocols communicate with the simulation engine by the exchange of actions.
These actions can be considered the basic building blocks of the MAC protocols. They include actions such as: sending packet, waiting a defined time the network to be free or counting down a defined time of free time slots. The use of these actions allows code reuse between the protocols and greatly reduces the complexity in the MAC protocols themselves.
It is not a simple task to implement a correct protocol. To simplify their verification protocols are designed as finite state machines. They are designed as mealy automata. As input they get information about the last action (packet dropped, collided, etc.) from the simulation engine. Based on this input and their internal state they chose their new internal state and return the next Action the protocol wants to be executed.
Protocols are instantiated by a traffic stream. Beeing on the MAC layer the details of the upper layers are hidden from the MAC protocols. They communicate with the traffic stram by a traffic source a class which provides simple packets. The protocols put new packets available in their input queue and transmit them after due time. The packet class encapsulates all the packet specific information, needed for the recorder to collect statistics.

4.2.3 Traffic streams

Traffic streams together with several other classes allow simulation of the upper layers of the network stack. Various transmission scenarios of unicast and multicast streams, using different transport protocols can be realized. Transmission streams are instantiated together with a recorder and a MAC protocol. The composition of big interchangeable customizable parts allows fine grained control of all parameters.

4.2.4 Recorder

The built in data collection has a number of data summarization primitives to assist the user in gathering network performance statistics during a simulation execution. The central package providing this functionality is the recorder. This design allows a recording facility to be added to every traffic stream which should be surveyed.
The various statistics tools are categorized into groups of annotated graphs. They produce standardized graphs like average graphs and packet graphs. This facilitates the task of the plotting tool to visualize the data.

4.3 Using *wireless Sim*

As shown, the simulator consists of a large number of Java objects which implement the behaviour of the network.
A user wanting to perform network simulations creates a Java main program instantiating the Java objets representing the various network elements comprising the simulation. After compilation, the execution of the newly created program allows the simulation defined by the composition of the objects.
In order to make batch processing possible it is desirable to be able to group multiple simulations. For this purpose the simulator introduce notion of lists of *simulation sets*. Multiple different simulations together with the style of comparison graphs wished can be combined to a simulation set and a list of them executed in batch mode.

Chapter 5

Performance evaluation of the enhanced MAC layer

5.1 Simulation parameters

In chapter 3, some enhancements to the MAC layer for multicast video transmission are proposed. In this chapter the performance of this modified MAC layer (*enhanced MAC layer*) for multicast transmission is compared to multicast transmission defined by the standard (*standard MAC layer*).

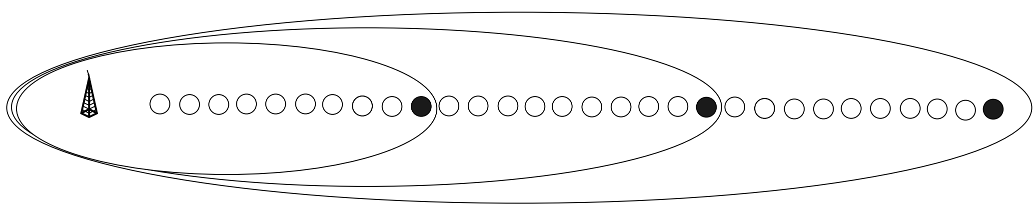

Figure 5.1: Receivers ordered by their distance from the base station; the stations responsible for sending acknowledgments are coloured in black

We have chosen the following parameters for the simulations:
Thirty receivers with a distance between 5 and 150 meters of the access point are subscribed to the video streams. The video is sent in three layers, one base layer and two optional enhancement layers. All 30 receivers, are subscribed to the base layer (layer 1). The 20 receivers with the best signal to noise rate (SNR) receive additionally the first enhancement layer (layer 2). Of this group the ten receivers with the best SNR receive additionally also the second enhancement layer (layer 3).
In each group receiving one layer, the receivers which have the worst channel conditions (measured as the mean SNR) are chosen to acknowledge the packets of this layer with an ACK. This means that senders 10, 20 and 30 are responsible to send ACKs for the packets received of layer one, two and three respectively (see figure 5.1)

All results are based on 120 seconds of video transmission.

5.2 Goodput gains of the enhanced Mac layer

We carried out the simulations with a bandwidth of 14, 200 and 300 kByte/s for layer 1 2 and 3, to demonstrate the advantages of the *enhanced MAC layer*

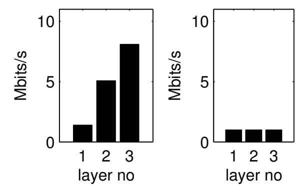

Figure 5.2: average transmission rate enhanced MAC layer (left) and standard MAC layer (right)

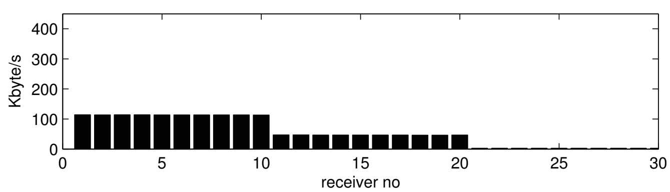

Figure 5.3: total goodput for each receiver in kbytes per second, standard MAC

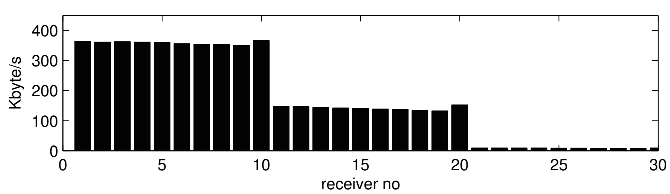

Figure 5.4: total goodput for each receiver in kbytes per second, enhanced MAC

The first observation in all simulations is that the physical transmission rate for the layers is chosen comparatively well (see 5.2) The base layer which has to be received by every station is sent at a mean transmission rate of 1.4 Mbit/s. The third layer only for the 10 closest receivers is sent at an average (arithmetic mean over all transmissions) of 8 Mbit/s.

The rate algorithm is controlled by the feedback of the ACKsenders. Different ACKsenders of the layers adapt the transmission rate to the multicast group of each layer.

This adaptation results in a goodput of the *enhanced MAC layer* up to four times as high compared to the *standard MAC layer* (5.3 and 5.4). This is remarkable because we have to keep in mind that transmission with

ACK and resending in case of packet loss adds up more to overhead for the *enhanced MAC layer*.

In these simulations the network is saturated by the video streams. This allows to visualize the goodput difference, but a bandwidth allocation which saturates the network is not realistic. The congestion results in a high packet loss rate. This would severely degrade video performance. But these simulations convey some important information. They show the bandwidth limits for both MAC layers. If the bandwidth of the layers is chosen below these limits, no packets are lost due to congestion. We see the bandwith available for video multicast is much higher in the *enhanced MAC layer*.

If lower values for the bandwidth of the layers are chosen like in the next simulations the more efficient transmission of the *enhanced MAC layer* results in a lower channel occupancy time (measured as percentage of the total time, where there is a packet transmission in the network) (see 5.8 in the next simulation).

5.3 Loss rate comparison

5.3.1 Error percentage in the absence of collisions

In the following simulations video was multicasted in three layers of 14, 28 and 50 KByte/s. With these parameters, congestion, which would make a comparison difficult[1] was not observed.

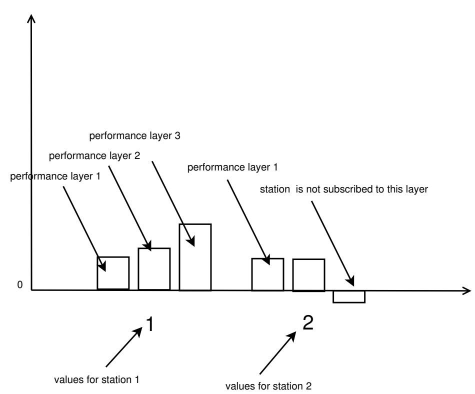

Figure 5.5: legend of the following figures

The multicasted video was the only data transmitted; no other traffic which could disturb the video transmission was sent. All packet losses observed here, are therefore packets lost, because of bad reception quality.

[1]The *standard MAC layer* suffers from congestion already at a lower bandwidth. There would be a lot more packets dropped in simulating the *standard MAC layer* . This would render a meaningful interpretation of the errors impossible.

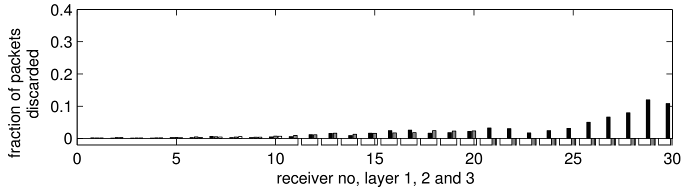

Figure 5.6: error rate standard MAC layer

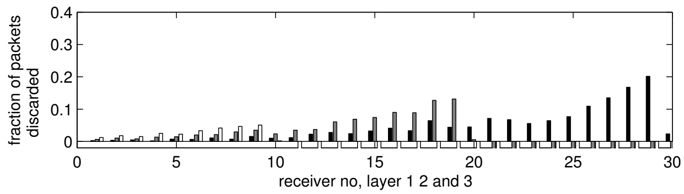

Figure 5.7: error rate enhanced MAC layer

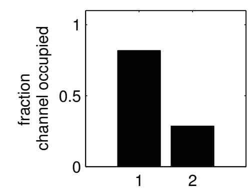

Figure 5.8: occupancy of the channel measured as the fraction of the complete time one station is sending, *standard MAC layer* (1) and *enhanced MAC layer* (2)

Figure 5.6 and 5.7 shows the error rate measured. The error rates for the *standard MAC layer* are low (up to 4% for the first 25 receivers) They grow with the decreasing reception quality of the receiver. The error rates of the *enhanced MAC layer* show a complicated behaviour: Compared to the *standard MAC layer* the error percentage is higher (error rates reach up to 12 % for the first 25 receivers). If we regard the different error rates of the layers for one station the error rate is increased for the higher layers. Generally we can say that with increasing distance of the station the error percentage of the losses is rising. An exception are the layers of the dedicated ACKsender (layer 3 for the station 10, layer 2 for the station 20, and layer 1 for the station 1). They have a very low error percentage.

It is obvious that the raw packet error rate before correction for the *enhanced MAC layer* is higher: The layers are transmitted at a higher rate. (see figure 5.2). Transmissions at a higher physical rate is more error prone[2]. While reducing the error percentage for the ACKsender itself, the figures show that retransmission based on the ACK reception of the dedicated ACK-sender does not reduce significantly the errors of the other stations subscribed

[2]see section 2.4.1 for details.

to this layer.
Figure 5.9 sheds light on this phenomenon.

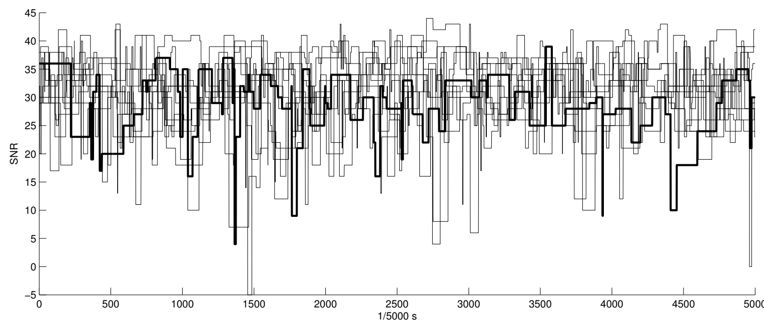

Figure 5.9: SNR varying with the time. One second of the SNR values of the ten receivers with the best average SNR is plotted. The SNR of the dedicated ACKsender is plotted in bold

Figure 5.9 plots the SNR values of the receivers with the best reception quality for a duration of one second (these receivers are all the receivers receiving the third layer). Based on the ACK of the tenth receiver whose SNR is plotted in bold the base station decides at which physical rate to send the layer and whether to resend a packet. Even if mean SNR values in one group are similar, momentarily they can differ quite a bit. The rate chosen by the algorithm based on the SNR of the dedicated receiver is at some instances to high for the others.

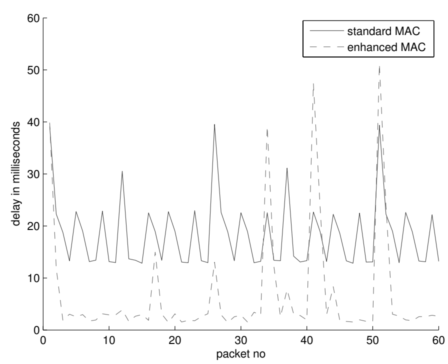

Figure 5.10: delays obtained for the *standard MAC layer* (1) and *enhanced MAC layer* (2) the delays of the first 60 packets are shown for packets of the first layer

Figure 5.10 shows the delays of the packets. Packets not received are plotted with a delay of 0 seconds. The delays stay always bounded below 1 sec. The delays measured for the *enhanced MAC layer* are below the delays of the standard multicast for most of the packets transmitted. There are some exceptions to the rule. When a transmission fails, packets are resent in the *enhanced MAC layer*, and the time of the backoff and retransmission accumulate to a higher delay.

5.3.2 Advantages of resending packets in a situation with high collisions

The same simulation was carried out, where we supposed that additionally to the video multicast, stations 1 to 15 send an UDP stream of 168 kbit/s: packet collisions frequently occur. Here retransmission based on the ACK reception of one representative of each group, helps reducing the errors.

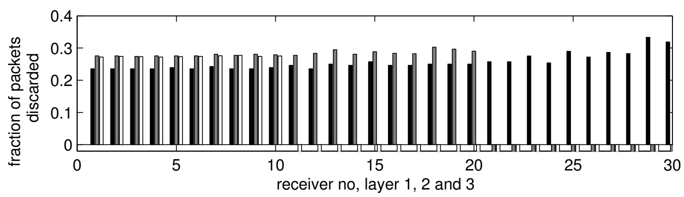

Figure 5.11: error rates *standard MAC layer* with. Station 1 to 15 are sending an UDP stream of 168 kbit/s

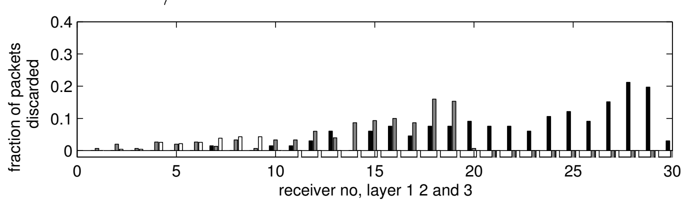

Figure 5.12: error rates *enhanced MAC layer* same setup with 15 background streams

The simulations of the *standard MAC layer* show a high error rate (see figure 5.11). As errors are not corrected by retransmission, every collision results in packet loss. Figure 5.12 shows that these additional errors can be almost completely eliminated by the *enhanced MAC layer*.

The remaining errors shown in figure 5.12 are mainly packets which are not received because of the bad reception quality of the receivers.

5.4 Conclusions

We can summarize the results as follows:

The difference in performance between the *standard MAC layer* and the *enhanced MAC layer* are based on two effects: retransmission and rate selection.

5.4.1 Retransmission effects

- Retransmitting packets based on ACK reception of a representative of each layer reduces the losses caused by collision significantly: Collision based losses are losses which are highly correlated between the receivers, therefore one dedicated receiver is able to "speak for the others" in his group.

- Resending packets based on the reception of the ACK of one dedicated receiver is not a good solution for the elimination of the packet loss based on corruption. While *average* reception qualities (and error rates) in one group are similar, the momentary SNR value differs distinctly. The SNR for one receiver can drop spontaneously for milliseconds because of deep fades of the media. These fades are at different time intervals for different receivers. The packets which are resent because the channel conditions of the one dedicated receiver are currently bad, are not correlated to the packets needed to be resent for others in the group.

5.4.2 Rate selection in multicast

- Rate selection is a critical point in reaching high throughput. With a wide gap in transmission speed (between 1 and 11 Mbit/s for 802.11b) which will grow if new standards emerge and are widely deployed (802.11a, 802.11g, 802.11n) this is a major point for improvement which can not be neglected. The proposed MAC layer with the CLARA Algorithm shows what performance improvements can be achieved.

- There are still some problems to solve, in adapting rate algorithms to multicast. The CLARA algorithm modified to work in multicast, chooses a rate for packets sent to a whole group of receivers based on the information included in the ACK of one dedicated receiver. CLARA is able to choose a transmission rate which is mostly correct. But the algorithm experiences some drawbacks in multicast.

- In unicast the algorithm is able to lower the rate for the duration of long fades, to a transmission rate, where packet transmission succeeds. In multicast, only the dedicated receiver benefits that the rate adaptation closely follows the fades. CLARA is able to chose a good average rate but it can not as closely model the instantaneous best rate for everybody in the group based on the reception of one receiver.

- The results show also that in multicast and especially in congested networks, even for the channel which sends the dedicated ACKs the rate selection works not as well as in the unicast case. This can be explained by the internal mode of operation of CLARA, The receiver transmits with each ACK (or CTS) information about the reception quality of the channel. The algorithm works well if it has recent channel information available. If packets are transmitted with a RTS/CTS sequence, the sender gets information about the state of the channel just before transmission of the packet. The *enhanced MAC layer* algorithm for multicast proposed in chapter 3, transmits the packets without RTS/CTS. The last information the sender receives to decide how the rate for the next data packet has to be chosen is how the the last data packet has been received. This information is sent to him via the reserved bits of the PLCP header of the ACK. As there is a backoff between the measurement of the channel and the new packet to send (which increases if collisions are frequent) the state of the channel is less accurately predicted.

The proposed MAC layer achieves significantly higher transmission. We observe at the same more packet errors. The error correction scheme is not the best scheme to correct errors and there may be better ones (see next chapter). But higher packet loss might be a characteristic to accept in multicast. The results show that there is still a considerable goodput gain even if we consider the additional error correction overhead on the application layer needed to eliminate these errors.

Chapter 6

Outlook

This work can be considered as pioneering work. Very few work has been done wireless LANs optimizing multicast video transmission. While it is challenging to tackle a new field, it has one drawback: the solutions obtained are not as polished they might be in an already well established field, where one can suggest another improvement to an already good solution.

Our proposed changes for the MAC layer show the performance advantages to be gained in multicast with a rate-selection optimized MAC layer. The change needed: Acknowledgement of packets is a minimal modification of the standard. Wireless stations implementing this are still compatible to stations incorporating the existing standard.

While we have shown the first steps for an optimized layered solution there are more questions to be answered than solutions proposed. To serve a starting point to future work we list some ideas to follow and some problems to tackle in this new area.

6.1 Modeling correlations in the channel fluctuations

The effects perturbing packet transmission in wireless LAN are well known; well founded channel models for one receiver exist. But studies of the long term *correlated loss behaviour* of multiple wireless stations are sparse.

Multipath fading is the effect which has the most impact on the short term channel variability. The correlation of multipath fading has been studied intensively because it is important for MIMO (multiple-input multiple-output) approaches a hot topic today: Fortunately for MIMO proposal and unfortunately for the design of a multicast MAC the fading experienced of multiple receivers appears to be uncorrelated.

The MIMO approach is based on this independence of the fading: In the wavelength wireless LAN is operating, an array of multiple antennas even if centimeters apart shows already uncorrelated fading effects. By choosing the antenna with the best instantaneous reception quality from the array of antennas it is estimated that throughput can be increased considerably.

For the evaluation of a clustering algorithm it will be important to model other effects which account for fluctuations on a longer time scale. There will certainly be correlation between receivers based on *shadowing effects.*

If one wants to study the performance of the SARC clustering algorithm, a good channel model that takes into account these long term correlation effects is a prerequisite.

As long as there are no good models the next step is to conduct real experiments to evaluate the proposed approach of adaptive clustering.

6.2 Inclusion of a hybrid FEC/ARQ error correction scheme in the MAC layer?

We have shown that retransmission schemes alone are not sufficient to correct uncorrelated errors in a group of receivers. A hybrid FEC/ARQ error correction scheme is more effective. One redundant FEC packet, can correct multiple independent losses of different receivers. With long fades causing successive packet errors FEC codes reach their correction limits. Therefore a hybrid solution[4] may be the best approach.

An interesting idea might be not to leave this task to the application layer but include a hybrid FEC/ARQ error correction mechanism in the MAC layer.

What are the advantages and drawbacks to include it here in the data link layer? First, it is the task of data link layer to provide reliable data transfer across the physical link. There are special mechanisms that provide reliability in unicast. It is natural to include a mechanism in multicast. Streaming solutions would be more independent from the special characteristics of the wireless LAN.

Second, an implementation at the MAC layer might be a lot more efficient than on the application layer: ACKs as feedback are far less overhead than acknowledgments on upper layers. Furthermore, the sender while sending a data packet reserves with the network allocation time the timeslot for the ACK. Feedback congestion can be avoided. At the MAC layer a solution can better differentiate between different causes of packet loss: collision, corruption, and congestion. Also it might be preferablye to include a rate selection algorithm into such a scheme. These advantages may justify the additional complexity of the MAC layer.

6.3 Standardized interface for cross layer communication

Our results showed that there are significant performance improvements to be gained in a jointly optimized architecture. Unfortunately today application designers of a video transmission streaming application have to treat the MAC and physical layer as a „black box“. A black box approach is great if the black box works well. But if it does not work, diagnosis is made difficult and there is little room for improvement. The high fluctuations of the channel characteristic to wireless LAN networks do not allow the MAC and physical layer to be a black box for application layers. There have to be some diagnosis tools provided for the application layer by the MAC layer, based on these it makes sense to control the parameters of the physical and MAC layer. The parameters of the MAC layer are application specific. For some applications, it is better to have a reliable transmission without errors. For other applications it is more important to send with a high bandwidth and some errors can be tolerated. In our approach for example enhancement layers are more tolerant to errors than the base layer.

It would be interesting to define a standardized communication interface between the layers to allow better error diagnosis and control of rate selection and error correction mechanisms. Today the wireless standards do not propose such an interface. With the next generation of WLAN standards on the horizon especially 802.11n we have to think about including such an interface. Such a solution still preserves advantages of a layered architecture but gives applications the performance possibilities of a customized solution.

Bibliography

[1] H. Schulzrinne, S. Casner, R Frederick, V. Jacobson, *A Transport Protocol for Real-Time Applications*, RFC 1889, January 1996

[2] Colin Perkins *RTP Audio and Video for the Internet* Addison Wesley, 2003.

[3] J. Zhou, H. Shao, C. Shen, M. Sun, *Multi-Path Transport of FGS Video* Packet Video (PV), April 2003

[4] A. Majumdar, D.G. Sachs, IV. Kozintsev, K. Ramchandran, M. Yeung *Multicast and unicast real-time video streaming over wireless LANs* in IEEE Transactions on Circuits and Systems for Video Technology, vol 12, no 6. June 2002

[5] J. Vieron, T. Turletti, K. Salamatian, C. Guillemot, *Source and channel adaptive rate control for multicast layered video transmission based on a clustering algorithm*, in EURASIP Journal on Applied Signal Processing, Vol 2004, No. 2, pp. 158-175, February 2004.

[6] IEEE 802.11, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications* Standard, IEEE August 1999.

[7] IEEE 802.11b, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, Higher-Speed Layer Extension in the 2.4 GHz Band* Standard, IEEE September 1999.

[8] H. Hashemi, *The Indoor Radio Propagation Channel* Proceedings of the IEEE, Vol. 81 No 7:943-986, July 1993.

[9] M. K. Simon and M. Alouini, *Digital Communication Over Fading Channels: A Unified Approach to Performance Analysis*, John Wiley & Sons, 2000.

[10] Gavin Holland, Nitin Vaidya, Paramvir Bahl, *A rate-adaptive MAC protocol for multi-Hop wireless networks*, Proceedings of ACM MobiCom, July 2001, Pages: 236 - 251.

[11] , Ad Kamerman, Leo Monteban, *WaveLAN-II: A high-performance wireless LAN for the unlicensed band* Bell Labs Technical Journal, Summer 1997, Vol. 2 Issue 3, pp. 118.

[12] C. Hoffmann, M. H. Manshaei, T. Turletti, *CLARA: Closed-Loop Adaptive Rate Allocation for IEEE 802.11 WirelessLANs,* WIRELESSCOM, June 2005 Maui, Hawaii.

[13] Kevin Fall *Network Emulation in the Vint/NS Simulator* Proceedings of the Fourth IEEE Symposiumon Computers and Communications, July 1999.