

Hybrid Rate Control for IEEE 802.11

Ivaylo Haratcherev

I.Haratcherev@ewi.tudelft.nl

Reginald Lagendijk

R.L.Lagendijk@ewi.tudelft.nl

Koen Langendoen

K.G.Langendoen@ewi.tudelft.nl

Henk Sips

H.J.Sips@ewi.tudelft.nl

Faculty of Electrical Engineering, Mathematics, and Computer Science
Delft University of Technology
Mekelweg 4, 2628 CD Delft, The Netherlands

ABSTRACT

Streaming multimedia content in real-time over a wireless link is a challenging task because of the rapid fluctuations in link conditions that can occur due to movement, interference, and so on. The popular IEEE 802.11 standard includes low-level tuning parameters like the transmission rate. Standard device drivers for today's wireless products are based on gathering statistics, and consequently, adapt rather slowly to changes in conditions. To meet the strict latency requirements of streaming applications, we designed and implemented an advanced control algorithm that uses signal-strength (SNR) information to achieve fast responses. Since SNR readings are quite noisy we do not use that information to directly control the rate setting, but rather as a safeguard limiting the range of feasible settings to choose from. We report on real-time experiments involving two laptops equipped with IEEE 802.11a wireless interface cards. The results show that using SNR information greatly enhances responsiveness in comparison to statistics-based rate controllers.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless Communication

General Terms

Experimentation, Measurement, Performance

Keywords

Rate control, MAC layer, SNR, Link adaptation

1. INTRODUCTION

It is anticipated that multimedia streaming over the Internet will have a significant share in tomorrow's communications. Also, end users increasingly seek mobility, thus

paving the way for extensive deployment of wireless technologies like IEEE 802.11. The joint effect is that support is needed for multimedia streaming over connections that include both fixed and wireless links. In this paper, we focus on the weakest part of such connections: streaming over a wireless link (the last hop). Such a link is the bottleneck for two reasons: First, communication over a wireless channel is simply not able to achieve the same quality (throughput, error rate, etc.) as its wired counterpart, which reduces the quality of the multimedia content that can be delivered. Second, in a mobile environment, the channel conditions can change rapidly due to changing distance between the stations (user mobility), Rayleigh fading, interference and so on. Since multimedia streaming applications must deliver their content in real time, they are very sensitive to jitter in packet delivery caused by retransmissions in the underlying transport protocols. Consequently, when using streaming applications, users experience reduced range compared to the case when less demanding applications like file downloading and web browsing are used.

With today's 802.11 products, the fundamental problems of wireless communication are aggravated by poor handling of the limited and imperfect resources (scarce spectrum, noisy medium) available to the radio. In particular, current transport protocols and device drivers do not actively control the user-available parameters of the 802.11 MAC layer; they use some default values instead. In this paper, we demonstrate that much can be gained by tuning the MAC parameters to the (fluctuating) channel conditions.

The IEEE 802.11 standard defines several MAC-level parameters (settings) that are available for tuning at the side of the Wireless Network Interface Card (NIC). The most important parameter available for adjustment is the transmit rate. In 802.11b, for example, it can be set to 1, 2, 5.5, and 10 Mb/s. Each rate corresponds to a different modulation scheme with its own trade-off between data throughput and distance between the stations. This can be seen in Figure 1, which was taken from [3]. It shows the simulated performance in terms of the throughput for each modulation scheme available in IEEE 802.11b versus the signal-to-noise ratio (SNR). Note that distance is related to SNR as $SNR \sim \frac{1}{dist^\alpha}$. More complex modulation schemes like CCK 11 offer a larger throughput, but also have increased sensitivity to channel noise, and thus provide a shorter operating range. Usually, one wants to extend the operating range as much as possible and, at the same time, to max-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiWac'04, October 1, 2004, Philadelphia, Pennsylvania, USA.
Copyright 2004 ACM 1-58113-920-0/04/0010 ...\$5.00.




Figure 1: Throughput vs. SNR for 802.11b modulation schemes.

imize the throughput. This can be done by proper (automatic) selection of the rate (modulation scheme) that gives the maximum throughput for certain conditions, for example, by selecting QPSK at an SNR of 5 dB (Figure 1).

In order to decide which rate is optimal at each specific moment, a control algorithm needs information about the current link conditions. Since it is difficult to measure link conditions directly, most rate-control algorithms use some form of statistics-based feedback, for example, user-level throughput (see Section 2). The main disadvantage of this indirect feedback is that it is inherently slow, causing communication drop-outs when the link conditions degrade rapidly (e.g., when the user moves fast). The short-term drop-outs are normally handled by frame retransmissions. This is acceptable for download applications whose main requirement is gross data throughput. The retransmissions, however, lead to a significant increase in (average) packet delay, and the variations in the number of retransmissions cause an increase in the jitter of the packet delay. Streaming applications are very sensitive to long packet delays and high jitter, and less sensitive to the overall throughput of the link (provided of course that this throughput is larger than the minimum throughput that the application requires). Consequently, streaming applications perform poorly under standard automatic rate control. In [4], for example, it is reported that switching from automatic rate control to a manually selected fixed rate extends the maximum distance between stations up to 40% for the flawless display of a video stream over an 802.11a wireless connection.

A logical way to cope with the slow accommodation characteristics of statistics-based feedback methods is to look for methods that use faster feedback, i.e., feedback that quickly provides up-to-date information about the channel status. Such a feedback - the SNR - has been theoretically discussed in previous works, but it has never been used in actual implementations. The main reason for this is that, in practice, it turns out to be very difficult to obtain a reliable estimate of the SNR (see Section 2). In this paper, we discuss an advanced approach to mitigate the SNR-related problems, and report on a practical implementation of this approach in a novel automatic rate-selection algorithm for IEEE 802.11a wireless connections.

The remainder of the paper is organized as follows. Section 2 gives a description of the existing rate control algorithms. Our improved solution is introduced in Section 3. The experimental results are discussed in Section 4. The conclusions and future plans are presented in Section 5.

2. CURRENT APPROACHES AND RELATED WORK

The IEEE 802.11 standard [7] and the 802.11a and 802.11b supplements [8, 9] do not specify an algorithm for automatic rate selection. So it are the companies manufacturing 802.11 interfaces, that have to come up with the algorithm: they have the freedom to design and implement their own proprietary schemes for control. To our knowledge, all of the known vendors of 802.11 equipment use statistics-based approaches for rate control. The control algorithms can be implemented in software as part of the device driver for the wireless NIC, but also as part of the chipset, which allows for precise control of individual frame (re-)transmissions. In the research community, another class of rate control algorithms has been studied. These control algorithms use SNR-related information as a feedback to improve the sensitivity to changes in link conditions. Up till now, however, such algorithms were not implemented in practical systems, and only simulation results were reported.

In the discussion below we will describe the main representatives of both classes (statistics- and SNR-based) of rate control algorithms. It should be noted that the performance of all algorithms can be improved by differentiating to packet length, because the probability of a packet being corrupted depends on the length of the transmission. Consequently, long packets should be transmitted at a lower rate than short packets.

2.1 Statistics-based automatic rate control

An easy way to obtain the necessary information on the link conditions is to maintain statistics about the transmitted data like the frame error rate (FER), acknowledged transmissions, and the achieved throughput. Since these statistics are directly related to the effective user-level data throughput, they inherently guarantee that this throughput is maximized on the long-term. These factors (simplicity and throughput) explain the dominance of statistics-based feedback in current 802.11 products.

2.1.1 Throughput-based rate control

In this approach, a constant small fraction (10%) of the data is sent at the two adjacent rates to the current one (an adjacent rate is the next higher or lower one available). Then, at the end of a specified decision window, the performance of all three rates is determined by dividing the number of bytes transmitted at each rate by their cumulative transmission times. Finally, a switch is made to the rate that provided the highest throughput during the decision window. Atheros uses this algorithm in the NIC driver that they provide for their 802.11a products based on the AR5000 chipset

To collect meaningful statistics, the decision window has to be quite large (i.e., about one second). On the one hand this makes the algorithm resilient to short-lived changes in the link quality caused by, for example, Rayleigh fading. On the other hand, it prevents swift reactions to long-lived

changes in link conditions, which noticeably affects the real-time performance of streaming applications.

2.1.2 FER-based rate control

In this approach, the Frame Error Rate (FER) of the data stream transmitted over the link is used to select an appropriate rate. The FER can easily be determined since under 802.11, all successfully received data frames are explicitly acknowledged by sending an ACK frame to the sender; hence, a missing ACK is a strong indication of a lost data frame. By counting the number of received ACK frames and the number of transmitted data frames during a rather short time window, the FER can be computed as the ratio of the two.

The FER can be used to select the rate setting for the next time window as follows [2]:

downscaling If the FER exceeds some threshold and the current rate is not the minimal rate, then switch to the next lower rate.

upscaling If the FER is close to zero (i.e., below a second threshold), probe the link at the adjacent higher rate with a few (usually even only 1) frames. If all of them get acknowledged, switch to that rate. To prevent the control algorithm from oscillating between two adjacent rates, the upscale action may be prohibited for some time after a downscale decision.

The width of the time window and the thresholds mentioned above are critical for the performance of the FER-based algorithm. The optimal settings of the parameters are dependent on the link and the application, but are generally fixed at design time. Again, this hampers the performance of streaming applications, since a time window tuned for quick responses of typical download applications (to changes in link conditions) yields unreliable FER statistics at low traffic rates. Hence, many frames are transmitted at a non-optimal rate.

2.1.3 Retry-based rate control

An improvement over the FER-based approach is to downscale immediately when the MAC is struggling to transmit a frame correctly over the link. That is, to select the next lower rate after a small number of unsuccessful retransmissions (usually 5-10 retries) [6, 12]. This approach is implemented in hardware, as precise control of the rate setting in between retransmissions (of the same frame) is required.

The advantage of the retry-based approach is that it combines a very short response time (a few frames) for handling deteriorating link conditions (downscaling) with a low sensitivity to traffic rates. The price to be paid is that the control algorithm is rather pessimistic. Relatively short error bursts cause long drops in throughput because upscaling to higher rates takes much longer than downscaling due to the need to collect a meaningful FER and to prevent oscillation.

2.2 SNR-based automatic rate control

A fundamental limit of indirect, statistics-based feedback is that it classifies link conditions as either "good" or "bad". This binary information provides some notion about the direction in which to adapt the rate setting, but does not suffice to select the appropriate rate at once. This leads to a slow step-by-step accommodation to large changes in con-

ditions, and introduces the risk of oscillation in stable conditions. A better approach is to use direct measurements of the link conditions.

The Signal-to-Noise Ratio (SNR) is directly related to the bit-error rate in the link and, hence, to the FER. Consequently, the SNR is linked to the packet delay and jitter, and the throughput, and holds the potential of providing rich feedback for automatic rate control [1]. Knowing the current SNR and the throughput-vs-SNR curves for each rate setting (e.g., Figure 1) solves the rate-selection problem instantly: we simply switch to the rate with the highest throughput for the current SNR. This can be implemented efficiently by means of a lookup table.

Despite the advantages, SNR-based rate control has not been applied in practice so far. This is mainly because of three reasons. First, in reality, for certain link conditions the relation between the optimal rate and SNR is highly variable. This is due to the imperfectness of the models describing the radio channel, and also because the link quality depends to some extent on other parameters as well. Second, it is not trivial to obtain a reliable estimate of the SNR of a link. Many radio interfaces provide only an uncalibrated Signal Strength Indication (SSI). Third, the rate controller, which is at the sending side, needs in fact the SNR observed at the receiving side. The problem of communicating the SNR information back is addressed in the emerging 802.11h standard [10], but the standard has not been finalized yet. In any case, this will be always an open issue for the products that do not support this standard. Also, being aimed at 5GHz, it is not possible for 802.11h to be a successor for the standards like 802.11b and 802.11g that operate at 2.4GHz. This would require defining a new supplement. Another drawback of 802.11h is that the SNR-related information (referred as Link Margin there) is transmitted back by the help of an additional management frame. The latter will increase the MAC overhead, and can cause the SNR information to be delayed due to contention for the medium, and thus not delivered on time.

Most work on using SNR information for automatic rate control is based on simulation and does not consider the practical difficulties of obtaining good SNR estimates. It concentrates on the way in which the noisy and drifting SNR (problem 1) can be used to determine the correct rate setting [1, 11]. Holland et al. [5] do address the issue of how to communicate back SNR values (problem 3), but their rate selection algorithm still relies on a straight SNR threshold technique. Another approach is discussed in [3], where the assumption is made that the channel is symmetric, meaning that the SNR observed at either station is very similar for any given point in time. This assumption allows Pavon et al. to use the SNR of the last ACK frame as an indicator of the SNR at the other side, and to use it for selecting the rate of the next data frame to be sent. Pavon et al. avoid the issue of estimating the true SNR out of an SSI reading (problem 2) by continuously adapting a table that maps SSI into throughput for the four 802.11b data rates. The adaptation is necessary to accommodate for varying noise levels. However, their control algorithm does not utilize the full potential of SNR information, since the adaptation only takes place on retransmissions; hence, adapting upwards is not supported well. Also, they provide simulation results only; therefore it is not clear how their algorithm will behave in practice. Finally, they do not discuss the dynamics of their

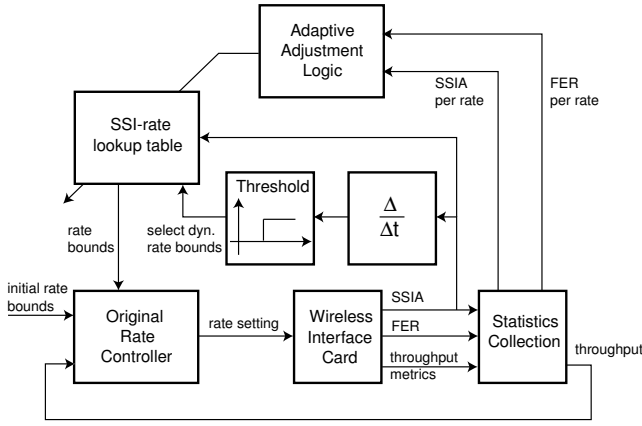


Figure 2: Structure of the hybrid rate controller.

approach, that is, how fast the rate setting reaches the optimal value for certain conditions, after these conditions were established.

3. HYBRID AUTOMATIC RATE CONTROL

In this section we describe an advanced rate control algorithm that combines the advantages of statistics-based and SNR-based methods. The goal is to support streaming applications by limiting the packet delay and jitter as much as possible, even at the expense of throughput when necessary. This requires a hybrid approach, since under stable link conditions we want to provide a robust, high-throughput connection (statistics-based approach), but under volatile conditions induced by, for example, user movement, we want to avoid retransmissions by rapidly switching to a lower rate (SNR-based approach). Another important goal is to design a working system, so we must address the practical problems associated with SNR-based methods as discussed in the previous section. Results of our experimental prototype will be presented in Section 4.

The core of our hybrid algorithm is a traditional statistics-based controller. More specifically, it is a throughput-based rate controller, which probes adjacent rates to determine if a rate switch is necessary (see Section 2.1.1). Figure 2 shows the complete structure of our hybrid rate controller; the throughput-based controller is part of the feedback loop depicted at the bottom of Figure 2, including the Wireless Interface Card and Statistics Collection component.

The decisions of the core controller can be overridden by a second feedback loop. This loop bounds the acceptable range of the Signal Strength Indication of the Acknowledged frames (SSIA) values for each rate, based on the specific knee in the Throughput-vs-SNR curve (cf. Figure 1). These SSIA bounds are implemented as a lookup table (SSI-rate lookup table component in Figure 2), which is indexed by the intended rate, given by the core controller. The bounds on the maximum rate are tightened when channel conditions change rapidly. The rationale of this scheme is that it is better to play safe in volatile conditions, and switch to a lower rate to minimize packet jitter by reducing the probability of a frame loss. Rapid changes in channel conditions are detected by sensing the changes of consecutive SSIA readings; the component is shown in the middle of Figure 2 as the

combination of a differential ($\frac{\Delta}{\Delta t}$) and a threshold detector. To account for the drift of the SSIA readings, we employ an adaptive adjustment logic that updates the values in the lookup table according to the recent history of channel conditions. This approach is similar to that described by Pavon et al. in [3], and may be regarded as an auto-calibration loop.

The operation of the hybrid control algorithm will be detailed further in two parts. First, we will describe the extended rate controller, which comprises all components in Figure 2 except the adaptive adjustment logic. Second, we will discuss the auto-calibration loop, that is, the details of the adaptive adjustment logic.

The SNR-based extension of the rate controller is centered around the SSI-rate lookup table. This table contains three SSI thresholds per rate: two for stable conditions (high and low thresholds), and an additional low threshold for volatile link conditions. These thresholds describe for each rate setting the SSIA values that are allowed for its use. The low threshold used in stable conditions is referred to as the *absolute* low SSIA value that provides acceptable performance (defined as FER < 10%) under ideal circumstances. The low threshold used under volatile conditions is referred to as the *dynamic* low SSIA value. The SSIA dynamics detector decides which of the two low thresholds (absolute or dynamic) will be used for bounding the rate selection for the next frame. The dynamics detector observes the differences between three consecutive SSIA readings. If both differences have the same sign, and both exceed a certain threshold, then the dynamic low threshold will be used, since conditions are changing fast and consistently. Otherwise, the absolute low threshold for stable conditions will be used.

The selection of the rate RS_{curr} of the next frame to be transmitted proceeds as follows:

```

once_per_decision_window() {
    RS_opt = find_opt_rate_by_stats();
    may_upscale = true;
}

for_each_packet() {
    RS_curr = probe_or_not(RS_opt);
    if (RS_curr > RUpBound)
        RS_curr = RUpBound;
    else if (RS_curr < RLoBound &&
            may_upscale) {
        RS_curr = RLoBound;
        try_upscale = true;
    }
    xmit_success = xmit_packet(RS_curr);
    if (try_upscale) {
        if (xmit_success)
            RS_opt = RS_curr;
        else
            may_upscale = false;
            try_upscale = false;
    }
}

```

Here, RS_{opt} is the optimal rate as calculated by the core statistics-based algorithm at the end of each decision window. The `probe_or_not` function determines whether or not a packet should be used for probing and, if so, returns one rate above or below RS_{opt} , otherwise returns RS_{opt} . The upper and the lower rate bounds ($RUpBound$ and $RLoBound$) are calculated from the SSIA thresholds in the lookup table ($SSIA_{tbl}$) and the current SSIA by the `calculate_bounds` procedure listed below. Note that when the selected rate

(RS_{curr}) falls below the lower bound ($RLoBound$), we attempt to upscale by transmitting a frame at the $RLoBound$ rate. If the transmission succeeds, we continue to use that rate. If it fails, we block additional attempts for the remainder of the decision window.

```

calculate_bounds() {
  for (i=max_rate_index, i > 0, i--) {
    if (curr_SSIA >=
        SSIA_tbl[i].lo_thld[is_dyn])
      break;
  }
  RUpBound = RateTable[i];
  for (i=0, i < max_rate_index, i++) {
    if (curr_SSIA <= SSIA_tbl[i].hi_thld)
      break;
  }
  RLoBound = RateTable[i];
}

```

The low SSIA thresholds determine the upper rate bound for the current SSIA value. Likewise, the high SSIA thresholds determine the lower rate bound. When there are rapid changes in channel conditions, the SSIA dynamics detector raises the `is_dyn` flag above, causing the algorithm to use the dynamic SSIA thresholds for calculating a stronger $RUpBound$.

The adjustment logic that updates the values in the SSIA lookup table at the end of each decision window operates as follows. First, an update is made for all rates that were used in the last decision window:

- The FER (in %) for each rate is calculated as:

$$FER = \frac{\#retries}{\#retries + \#good_packets} * 100$$

- If the FER is larger than 5%, then the low SSIA thresholds are updated by the formula:

$$lo_thld_{new} = lo_thld_{old} + (FER - 5) * a + b$$

- Otherwise, if the FER is less than 1%, and the observed average of the SSIA for the rate is less than lo_thld_{old} , then the SSIA thresholds are updated by:

$$lo_thld_{new} = lo_thld_{old} - (1 - FER) * c + d$$

- Finally, the dynamic and the high SSIA thresholds are updated:

$$hi_thld = lo_thld_{new} + e$$

$$lo_thld_{dyn} = lo_thld_{new} + f$$

Here, a , b , c , and d are coefficients that determine the accommodation properties of the adaptation algorithm. In the current implementation, e and f are constants, but in the future, they will represent functions. The e function accounts for the magnitude of the drift in the SSIA-FER relations. The f function determines how pessimistic the algorithm will be in the case of rapidly changing channel conditions, thus influencing the responsiveness of the rate controller. Currently, the values used are $a = 0.1$, $b = 1$, $c = 3$, $d = 1$, $e = 10$ and $f = 5$. However, they are subject to additional tuning, depending on the model of the controlled object (in our case the wireless interface). Therefore, as our model may be refined in the future, the coefficients can take different values.




Figure 3: Resolving specific cases for SSIA thresholds update.

After updating the thresholds for the rates used during the decision window, the thresholds for the unused rates may need to be updated as well to preserve the monotonic value growth in the lookup table. An example of a specific update situation is shown in Figure 3, where the thresholds for rate indexes 1 and 4 have already been updated according to the procedure outlined above. Next, the thresholds for the other rate indexes are checked. Rate index 2 is found to violate the monotonicity requirement, and its value is increased to that of the preceding rate index 1. No further adjustments are needed. To handle potential conflicts between threshold updates, the adjustments are performed from the lowest rate to the highest rate.

4. PERFORMANCE EVALUATION

To demonstrate that it is possible to design a practical SNR-based automatic rate controller, we have implemented our hybrid algorithm on real hardware and carried out a number of experiments executed in real-time.

4.1 Experimental setup

Our experimental setup consists of two laptops, equipped with a Proxim Skyline 802.11a wireless network interface card supporting eight different data rates, ranging from 6 to 54 MB/s. These cards are based on the Atheros 802.11a AR5000 chipset. The advantage of using this chipset is that it only implements the lowest MAC functionality (e.g., basic medium access, retransmissions control, etc.) and leaves the remainder to the device driver. This provides us with the flexibility to modify the original MAC protocol, and even to replace the automatic rate controller with our own hybrid approach.

The laptops are running Linux, and we developed the AR5000 device driver ourselves since only a Windows version was supported by Atheros. The device driver includes our hybrid automatic rate controller discussed in the previous section. For experimentation purposes the driver can be instructed to disable the SNR-based components so that we could measure the performance of the original throughput-based controller developed by Atheros. We modified the popular Netperf 2.2 tool to set the low-level MAC parameters exported by our device driver: maximum number of retransmissions, rate setting, Tx power, etc. We also used Netperf to generate data traffic and to collect statistics about

the raw performance of the wireless connection between the two laptops in a standard office environment (with concrete walls).

4.2 Results

To obtain basic insight into the performance of the wireless network cards, we conducted a first experiment in which we placed one laptop at different positions in the corridor. At each position we exercised all eight 802.11a rate settings logging the throughput reported by Netperf and the signal strength (of the acknowledgements) as reported by the network interface. Figure 4 shows the relationship between signal strength, rate setting (modulation scheme), and throughput for our experimental setup. In addition to the raw measurements, Figure 4 includes a smoothed¹ curve for each rate setting to better visualize the cliff behavior. The infeasibility of directly implementing a SNR-based rate control algorithm can be seen for example by looking at the results for 24 Mbit/s. Note that the data points represent averaged values over time. But still, some of the points (like those at a SSIA of 23 with throughput below 2Mbits/s) suggest that the cliff for that rate should be at about a SSIA of 23, while others on the left determine the switch over point to be around 15. Thus, straightforward SNR-based control will lead to unstable, and sometimes erratic behaviour.

The SSI readings that we obtained from the chipset ranged in practice from 0 (complete loss of connection) to about 70 (cards next to each other). Unfortunately, the documentation about the AR5000 chipset does not specify the unit of these measurements. Therefore we performed a series of open-space experiments in which we varied the distance, which revealed that the obtained SSI values are in dB over some reference point. This reference point is most likely linked to the noise floor of the receiver, thus we end up getting 'SNR-related' information. Our experiments show that the noise floor of the receiver is recalibrated periodically, which contributes to the drifting in time of the throughput-SSI curves. It is important to note here that, since we use a self-adapting algorithm, the reference point of the SSI readings does not influence the operation of the algorithm. This is provided, of course, that the reference point does not change too fast (within a second), which is the case in our experimental setup.

Figure 4 clearly shows that in reality, the obtained SSI feedback is quite noisy, especially around the edge of each cliff. In addition, the noise level for more complex modulation schemes (i.e., higher data rates) is quite high, even on the flat part of the curve. For example, at an SSI of about 42, the throughput of the three top rates suddenly plummets to a rate of just 1 Mbit/s. This is most probably due to complete Rayleigh fading of some of the OFDM carriers used in the 802.11a radio at that particular position in the corridor.

A second observation that can be made from the curves in Figure 4 is that, in our experimental setup, it does not make sense to use the 9 Mbit/s rate setting at all, since this setting does not outperform any other rate setting for any SSIA value. Likewise, the usability of the 36 Mbit/s rate setting can be questioned.

¹To smooth the data, we first excluded values that violate the monotonic increase of the curves. Then, we interpolated the data to obtain regularly spaced values. Finally, we ran the data through a low-pass FIR filter.




Figure 4: Throughput vs. SSI of acknowledgments (SSIA).

Finally we like to point out the throughput reductions caused by the transmission of the MAC headers at a fixed (low) rate. At the 6 Mbit/s setting, we logged a Netperf throughput of 4.9 Mbit/s, hence, the MAC overhead is limited to 18%. At the 54 Mbit/s setting, the Netperf throughput is limited to just 23.8 Mbit/s, hence, the overhead has grown to 56%. Note that the actual throughput, and consequently the MAC overhead, depends on the frame size used. In our case, the packets are 1024 bytes long, and we did not use packet fragmentation, so the frame size is equal to the packet size.

We conducted a second experiment to verify the assumption that the communication channel is symmetric, so that we can use the signal strength of the acknowledgements (SSIA) observed at the sender instead of relaying the true signal strength (RxSSI) observed at the receiver. The experiment involved a person walking around in a random fashion with one of the laptops (the other laptop remained stationary), and recording both the SSIA at the transmitting side and the SSI at the receiving side for each packet. Figure 5 shows how the signal strength observed at the receiver changed during the 140 sec experiment (packets are streamed at a rate of 100 per second). Note that during stable conditions (standing still), the noise in the SSI is limited to ± 2 dB; while moving, however, the fluctuations are much larger (up to ± 17 dB).

Figure 6 provides a histogram of the differences between the signal strength observed at the sender (SSIA) and receiver (RxSSI). Even though the SSIA and RxSSI readings are not captured from the channel at the same time (the ACK follows the DATA packet), it is quite clear that the SSIA and RxSSI values are strongly correlated. For 74% of the packets, the difference falls within the stationary ± 2 dB noise range, and even the largest differences are in line with the noise during movement. Therefore, from a control perspective, we can safely use SSIA instead of RxSSI, thus avoiding the problem of relaying the SSI at the receiver back to the rate controller operating at the sender.




Figure 5: SSI at the receiving side (RxSSI), random walk.




Figure 6: Histogram of the difference SSIA - RxSSI.

In the third experiment we compared the throughput-based rate controller supplied by the manufacturer of the chipset (Atheros) with our hybrid controller. To allow for a head-to-head comparison under the same conditions, we decided not to use a real-life scenario with a walking person, since repeating the experiment (i.e., regenerating the exact same channel conditions) would be close to impossible. Instead, we measure the Step Response Function (SRF) in an artificial setting to determine the responsiveness, robustness, etc. of the rate control algorithms.

In general the SRF is used to study how some control system reacts when a certain input parameter (in our case, the SSI) changes instantly, that is, in one step from one value to another. In our SRF experiment we placed one of the laptops in a microwave oven, started streaming data to the other laptop, and then closed the door rapidly, causing the SSI to swiftly drop from about 35 (good) to about 10 (bad), because the shielding of the microwave oven blocks almost all radio waves in the 5 GHz band. After about 3 seconds the door was opened again, generating a second step in the channel conditions (from bad to good). Since the

microwave door was opened and closed manually, the exact channel conditions (i.e, SSI readings) show small variations between experiments (cf. figures 7 and 8). Nevertheless, we can reproduce the SSI steps with enough accuracy to show the (large) differences between the throughput-based and hybrid rate control algorithms. To stress these differences we impose a traffic load that corresponds to a demanding streaming-video application; Netperf was instructed to send 1024 bytes long packets at a rate of 100 packets/s, resulting in an effective data rate of 800 Kbit/s.

Figures 7, 9, and 11 show the SSIA readings (input), generated rate settings (output), and induced packet latencies (performance), respectively for the original throughput-based rate control algorithm. Figures 8, 10, and 12 show the same for our hybrid algorithm. According to expectations the throughput-based controller cannot accommodate fast enough to the step down in channel conditions (around packet number 1000) because of the 1 second decision window. When it does respond (around packet number 1080), the controller decides to step down one rate setting (from 54 to 48 Mbit/s), which is too little too late. All packets transmitted during the closing of the microwave get lost; the latencies for these packets are labeled inf(inity) in Figure 11. As a consequence, no acknowledgements are received, causing a gap in the SSIA readings shown in Figure 7 and, hence, a lack of throughput statistics leading the controller to maintain the current (way too high) setting. This situation continues until the channel conditions improve after opening the microwave door, and the controller can finally decide to step down further (from 48 to 36 Mbit/s) around packet number 1380. Our hybrid controller, in contrast, responds almost immediately to the drop in channel conditions and overrules the decisions of its internal throughput-based controller based on the SSIA thresholds. Consequently, almost all packets are transmitted at the appropriate rate and service continues smoothly. When the conditions improve (around packet number 1290) the hybrid controller is also quick to respond and switches to higher rate settings.

The step-response function in Figure 9 shows the workings of the throughput-based controller in quite some detail. First of all, we can clearly see the 1 second decision window in the rate selections, for example, packets numbered 1380 to 1680 are grouped into three windows with rates 36, 24, and 36 Mbit/s respectively. Also, the probing at adjacent levels (1 in 10 packets) is visible as the light bars up and/or below the dark bar of the selected rate. Figure 10 shows that the hybrid algorithm does probe at adjacent levels too, but that quite often the SSIA-based thresholds overrule an up-probe. As an example, consider the time frame from packet number 1500 to 1750. In this period, the majority of the packets is sent at the 48 Mbit/s rate, some are sent at the adjacent lower rate (36 Mbit/s), but none at the higher rate (54 Mbit/s). This behavior follows from the data in Figure 4, which shows that around 35 dB the 54 Mbit/s throughput curve is still rising, signalling unstable conditions.

The packet latency, as reported in figures 11 and 12, is measured as the time between the moment the device driver inserts a packet in the send queue, and the moment that the chipset starts the actual, physical transmission. The latter time stamp is returned (by means of an interrupt) to the device driver when the acknowledgement packet is received. The chipset automatically takes care of retransmissions. If the maximum number of retries (10 in our case) does not re-




Figure 7: SSIA for the statistics-based algorithm.




Figure 8: SSIA for the hybrid algorithm.




Figure 9: Step response function - statistics-based algorithm.




Figure 10: Step response function - hybrid algorithm.




Figure 11: Packet latency - statistics-based algorithm.




Figure 12: Packet latency - hybrid algorithm.

sult in a successful transmission, the packet is dropped, and an interrupt is generated to inform the driver. The device driver then marks the latency for the dropped packet as infinite. Under normal conditions the send queue is (almost) empty, but when retransmissions are issued, the number of packets in the queue grows. When the control algorithm then switches to the “right” rate, the packets in the queue are transmitted back to back until the queue becomes empty again. This effect is clearly visible in the latency graphs: after a sequence of dropped packets, latency jumps and then rapidly declines to (near) zero. Streaming applications are quite sensitive to latency and usually put a strict requirement on the maximum latency for any packet.

The bottom line is that, from the point of view of a streaming application, the difference between the two automatic rate controllers is huge. With the original throughput-based controller, the stream is halted for about 2 seconds (205 packets are lost), causing a complete outage in service to the user. With the hybrid controller, the user may experience some service degradation depending on the application’s ability to cover up the few (5) dropped packets. In the case of a streaming video application, the difference is between losing connection and observing some artifacts in a few frames; a significant difference indeed.

5. CONCLUSIONS AND FUTURE WORK

In this paper we addressed the problem of supporting real-time streaming applications over a wireless link. The strict latency requirements of streaming applications and the fluctuating link conditions require careful handling in all layers of the protocol stack. At the lowest level, the IEEE 802.11 standard includes several tuning parameters, of which the transmission rate (selected per packet) is the most important one. Standard device drivers for today’s wireless 802.11 products are based on gathering statistics and, consequently, adapt rather slowly to changes in conditions. Using signal strength (SNR) information is appealing because it provides instant information about the channel conditions. However, since SNR readings are quite noisy, have a wide variance, and are unstable, SNR cannot be used in practice to directly select the rate.

We have designed and implemented a hybrid rate controller that uses the available SNR information as a safeguard. This hybrid controller is built around a traditional statistics-based rate controller. The SNR information is used to limit the range of feasible settings from which the core controller can choose. This prevents the transmission of packets at a too high rate when channel conditions suddenly change for the worse. In a controlled experiment, we measured the step response function of both the statistics-based rate controller provided by the manufacturer (Atheros) of the 802.11 chipset that we employ, and that of our hybrid rate controller. The performance difference is large: whereas the statistics-based algorithm incurs a total dropout due to its slow response, the hybrid algorithm adjusts almost instantly and loses only a few packets. We conclude that SNR information can be used to the advantage despite the practical problems of the noise and drift usually associated with it.

As future work we plan to investigate the benefits of including other parameters to be controlled (like packet fragmentation and Tx power). We will also study the possible use of inter-layer QoS negotiations, so that the application

will be able to influence the algorithm controlling the radio parameters. And last but not least: more-complex real scenarios with multiple nodes will be researched to improve our approach.

6. ACKNOWLEDGMENTS

We thank Tom Parker for commenting on the draft version of this paper.

7. REFERENCES

- [1] K. Balachandran, S.R. Kadaba, and S. Nanda. Channel quality estimation and rate adaptation for cellular mobile radio. *IEEE Journal on Selected Areas in Communications*, 17(7):1244–1256, July 1999.
- [2] B.E. Braswell and J.C. McEachen. Modeling Data Rate Agility in the IEEE 802.11a WLAN Protocol. In *OPNETWORK 2001*, March 2001.
- [3] Javier del Prado 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, Anchorage, Alaska, USA, May 2003.
- [4] I. Haratcherev, K. Langendoen, I. Lagendijk, and H. Sips. Application-Directed Automatic 802.11 Rate Control: Design Rationale and Preliminary Results. In *ASCI 2003*, pages 56–60, Heijen, The Netherlands, June 2003.
- [5] Gavin Holland, Nitin H. Vaidya, and Paramvir Bahl. A rate-adaptive MAC protocol for multi-hop wireless networks. In *ACM MOBICOM'01*, Rome, Italy, July 2001.
- [6] A. Kamerman and L. Monteban. WaveLAN-II: A High-Performance Wireless LAN for the Unlicensed Band. *Bell Labs Technical Journal*, pages 118–133, Summer 1997.
- [7] IEEE standard 802.11. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 1999.
- [8] IEEE standard 802.11a supplement. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. High-speed Physical Layer in the 5 GHz Band, 1999.
- [9] IEEE standard 802.11b supplement. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Higher-speed Physical Layer in the 2.4 GHz Band, 1999.
- [10] IEEE standard 802.11h supplement. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Spectrum and Transmit Power Management Extensions in the 5GHz band in Europe, 2003.
- [11] T. Ue, S. Sampei, N. Morinaga, and K. Hamaguchi. Symbol rate and modulation level-controlled adaptive modulation/TDMA/TDD system for high-bit-rate wireless data transmission. *IEEE Transactions on Vehicular Technology*, 47(4):1134–1147, November 1998.
- [12] A.J. van der Vegt. Auto Rate Fallback Algorithm for the IEEE 802.11a Standard. Technical report, Utrecht University, 2002.