RZ 3572 (# 99582) 01/03/05
Electrical Engineering 33 pages

Research Report

Dynamic Data Rate and Transmit Power Adjustment in IEEE 802.11 Wireless LANs

Pierre Chevillat, Jens Jelitto and Hong Linh Truong

IBM Research GmbH
Zurich Research Laboratory
8803 Rüschlikon
Switzerland

Research
Almaden · Austin · Beijing · Delhi · Haifa · T.J. Watson · Tokyo · Zurich

Dynamic Data Rate and Transmit Power Adjustment in IEEE 802.11 Wireless LANs

Pierre Chevillat, Jens Jelitto, Hong Linh Truong, IBM Zurich Research Laboratory

Abstract

In this paper a novel link adaptation algorithm is proposed that is capable of adjusting the transmit power level and the data rate jointly to the radio channel conditions. The proposed method relies solely on link quality information available at the transmitter by employing the reception or non-reception of the acknowledgment frames as a measure of the channel quality with respect to the power level and data rate. The method is fully compatible with the 802.11 wireless LAN standard. In contrast to many other proposals, it neither relies on the RTS/CTS protocol nor requires a feedback channel to transmit link-quality estimates from the receiver to the transmitter. Different strategies for optimizing data rate and power level are given. These depend on the scenarios considered, the number of active stations, and the service requirements. The two main strategies are either to drive the system towards the highest possible data rate and adjust the rate and power levels accordingly ("high-performance" mode) or to focus on power saving, possibly trading this for other performance criteria such as throughput or delay performance ("low-power" mode). Other special cases, such as power or rate only adaptation, are also discussed. It can be shown that in most cases the best choice for achieving low transfer times, maximizing throughput, and alleviating the hidden terminal problem is to transmit at the highest possible rates and with high power levels. This "high-performance" mode of operation also minimizes the transmission times, which in turn maximizes the time for putting idling components into a sleep mode, thereby minimizing overall power consumption.

Key words

802.11 Wireless LANs (WLANs), dynamic link adaptation, power and rate adaptation, throughput optimization, power efficiency

1 Introduction

Broadband wireless LANs (WLANs) based on the IEEE 802.11a [1] and 802.11g [2] physical layer standards support multiple data rates, which enables wireless stations (WSTAs) to select the appropriate transmission rate depending on the required quality of service and the radio channel conditions. To maximize the throughput or minimize the transmission delay, using a high data rate is the right choice in many cases. However, if the distance to the receiver is long, using a high data rate may lead to an excessive number of retransmissions owing to the low signal-to-noise ratio (SNR). This may result in performance degradation or even in a total loss of communication. In such a situation, a more robust but lower data rate may achieve a better throughput. Other parameters that affect overall system performance need also be considered. For battery-powered devices such as laptops and PDAs, transmission power awareness is crucial to save energy and prolong battery life. Reducing the transmit power level, in particular when the distance to the receiver is small, also helps reduce interference with neighboring WLANs. To achieve the highest possible system performance, it is therefore mandatory to use an automatic link adaptation algorithm that allows a station to adapt its transmission parameters to the actual conditions of the wireless channel.

The interest in designing a link adaptation mechanism for 802.11 WLANs is large because the basic 802.11 MAC standard [3] does not specify any procedure for rate switching or power level selection. Many papers have appeared in this area, most of them dealt either with rate adaptation [4] - [11] or with transmit power control [12] - [15] only. A few papers described algorithms that combine rate adaptation with power control [16] - [18].

The classical way for performing link adaptation is to rely on feedback from the receiver. In this approach the channel quality is estimated from SNR, received signal strength, or packet error rate measurements, and the transmit rate or power level to be used in future transmissions is derived. This information is then sent back to the transmitter over a feedback channel [19, 20]. Unfortunately, the 802.11 MAC standard does not provide any protocol means for the receiver to inform the transmitter about the link quality or the transmission rate to be used. Link adaptation methods relying on a feedback channel can therefore not be employed or require a modification of the 802.11 standard.

A number of proposed algorithms make use of the so-called “RTS/CTS” protocol to exchange the information needed for the adaptation. The RTS/CTS protocol is defined in the 802.11 standard to combat the well-known “hidden” terminal problem. Before sending a data frame a transmitting station makes a reservation for the wireless channel by sending a short “Ready-To-Send” (RTS) frame. The receiving station replies with a “Clear-To-Send” (CTS) frame. The transmitting station proceeds with the transmission of the data frame after having received the CTS frame. All other stations in the vicinity of the transmitter and the receiver that have received these two frames will defer their own transmissions until the end of the reservation. Thus, the RTS/CTS protocol can be exploited for the exchange of link adaptation parameters such as rate and packet size [4], or interference margins and transmit power levels [12].

An interesting idea is described in [9], in which the access point (AP) broadcasts special beacon frames that allow the wireless stations (WSTAs) to detect their location relative to the AP and to determine the transmit rate to be used.

Another alternative is to measure the received signal strength and derive the transmit parameters from that measurement [11, 21]. The transmit power level can also be included into the packets, such as in the method described in [17], in which a combined rate and power adaptation for 802.11a WLANs operating under the Point Coordination Function (PCF) is described. The main idea is to include the information about the transmit power level in the SERVICE field of a MAC frame which is sent by the Point Controller (PC). This information will allow the receiver of the frame to estimate the path loss between itself and the PC, and to determine the best rate-power pair to be used for future transmissions. A similar method is defined in [18] for the 802.11 Distributed Coordination Function (DCF). It uses the RTS/CTS frames sent at full power prior to any data transmission to estimate the path loss. The WSTAs determine the best power-rate combination from a table indexed by the packet length, the path loss condition, and the retry counters of the DCF procedure. The table entries are calculated off-line.

All algorithms described so far require either a modification of the 802.11 standard or the co-operation of the peer station, e.g. when making use of the RTS/CTS procedure. These drawbacks can be avoided by employing only information available at the transmitter. In [5], the transmitter switches between two fixed data rates, with the higher rate as the default operating mode. After two consecutive transmission errors it uses the lower rate, and returns to the high one after ten successful transmissions or after a time out. This technique was enhanced in [10] to handle multiple transmission rates and varying channel conditions. In [22] the mechanism of [10] is employed to adapt the transmit power levels to the channel conditions; however, the transmission rate is maintained at a fixed value.

In this paper we extend the ideas proposed in [10] and [22] to a new link adaptation algorithm that is capable of adjusting jointly the transmit power level and the data rate to the radio channel conditions. The method is fully compatible with the 802.11 wireless LAN standard. It relies neither on the RTS/CTS protocol nor requires a feedback channel to transmit link-quality estimates from the receiver to the transmitter. Instead, the proposed scheme relies solely on information available at the transmitter by using the ACK frames.

In the next section we will discuss some fundamental link adaptation issues. In Section 3 the proposed joint power and rate adaptation algorithm is presented. A comprehensive evaluation of the performance of the proposed scheme is given in Section 4. Finally, Section 5 concludes the paper.

2 Link adaptation principles

Our goal is to optimize throughput and transmit power for a given channel, i.e., to transmit as many bits with as little energy as possible. The optimization strategy depends on several factors, such as the service requirements (e.g. target data rate, delay constraints), power consumption, and battery lifetime, but also on coexistence and fairness considerations in a wireless network.

The parameters that can be controlled by a link adaptation algorithm are the transmit power P_{D}, data rate r_{D}, and packet length l_{D}. Although the throughput depends also on the packet length [23], it has been shown in [22] that fragmenting large packets as defined in the 802.11 MAC standard [3] is not an appropriate means for increasing throughput because every fragment is acknowledged separately. Hence, the overhead increases linearly with the number of fragments. For this reason, only transmit power and data rate adaptation will be considered in this paper. Let us define the effective throughput of a WLAN as

$$TP = \frac{N_{\mathrm{b}}}{T}, \tag{1}$$

where N_{b} is the total number of successfully transmitted data bits in the observed time interval T. The time interval T includes the transmit times T_{D} of the data packets and overhead times T_{OH} such as $DIFS$ ("Distributed Interframe Space"), $SIFS$ ("Short Interframe Space"), acknowledgment times, backoff times and other idle times (e.g. waiting times for new packets in a system that is not in saturation). In addition, let us define the average transmit power as

$$\bar{P}_{\mathrm{tx}} = \frac{\sum_{k=0}^{K-1} P_{\mathrm{D,k}} T_{\mathrm{D,k}} + \sum_{l=0}^{L-1} P_{\mathrm{OH,l}} T_{\mathrm{OH,l}}}{T}, \tag{2}$$

where $P_{\mathrm{D,k}}$ is the transmit power of the k-th data frame transmitted, $T_{\mathrm{D,k}}$ the transmit time of the k-th data frame, and K the total number of data frames transmitted, including erroneous frames. $P_{\mathrm{OH,l}}$ is the power associated with the l-th overhead time $T_{\mathrm{OH,l}}$ (e.g. to transmit an ACK frame, or the power consumed during idle times $DIFS$, $SIFS$), and L the total number of overhead events. In the simulation results presented in the next section, the power consumed during idle times (e.g. $DIFS$, $SIFS$) is not considered and set to zero. Hence, as long as no RTS/CTS mechanism is enabled in the system, the overhead energy consumption is determined by the transmission power P_{A} and transmission time T_{A} of the ACK frames to acknowledge successful data transmissions. Note that even in this case, K and L are not equal because of unsuccessful transmission attempts. From these two definitions the power efficiency can be derived as

$$\eta_{\mathrm{TP}} = \frac{TP}{P_{\mathrm{tx}}} = \frac{N_{\mathrm{b}}}{\sum_{k=0}^{K-1} P_{\mathrm{D,k}} T_{\mathrm{D,k}} + \sum_{l=0}^{L-1} P_{\mathrm{OH,l}} T_{\mathrm{OH,l}}}. \tag{3}$$

The denominator describes the overall transmit energy consumption associated with the successful delivery of N_{b} data bits. To maximize the power efficiency, the times T_{D} and T_{OH} and the powers $P_{\mathrm{D}}$and P_{OH} should be kept as small as possible for the given service requirements.

The transmission times $T_{\mathrm{D}}(l_{\mathrm{D}}, r_{\mathrm{D}})$ and $T_{\mathrm{A}}(l_{\mathrm{A}}, r_{\mathrm{A}})$ depend on the frame lengths l_{D} and l_{A} and the transmission rates $r_{\mathrm{D}}$and r_{A} of the data and ACK frames, respectively. For 802.11a, the time to transmit a data frame of l_{D} bytes at rate $r_{\mathrm{D}}$is [1]:

$$T_{\mathrm{D}}(l_{\mathrm{D}}, r_{\mathrm{D}}) = T_{\mathrm{PA}} + T_{\mathrm{SIG}} + T_{\mathrm{SYM}} \cdot \left\lceil \frac{16 + 8 \cdot l_{\mathrm{D}} + 6}{N_{\mathrm{DBPS}}(r_{\mathrm{D}})} \right\rceil, \tag{4}$$

where T_{PA}, T_{SIG} and T_{SYM} are the times to transmit the preamble, the signal field and one OFDM symbol of the 802.11a OFDM PHY layer, respectively. N_{DBPS} is the number of data bits encoded within one OFDM symbol and depends on r_{D}. The operator $\lceil\rceil$ returns the smallest integer value greater than or equal to its argument value. In 802.11a, these parameters have the values $T_{\mathrm{PA}} = 16\,\mu$s, $T_{\mathrm{SIG}} = T_{\mathrm{SYM}} = 4\,\mu$s, $N_{\mathrm{DBPS}} = 4 \cdot r_{\mathrm{D}}$. Possible transmission rates are r_{D}= 6, 9, 12, 18, 24, 36, 48, and 54 Mbps. If we assume a packet length of $l_{\mathrm{D}} = 1000$ bytes, the transmission times for the different data rates are $T_{\mathrm{D}}(l_{\mathrm{D}} = 1000\,\text{bytes}, r_{\mathrm{D}} = \{6, 9, 12, 18, 24, 36, 48, 54\}\text{Mbps}) = \{1360, 912, 692, 468, 356, 244, 188, 172\}\mu$s. Hence, for $l_{\mathrm{D}} = 1000$ bytes, the transmission time at a rate of 54 Mbps is approximately 1/8 of the time required at a rate of 6 Mbps.

The necessary transmit power for a certain service requirement depends on the transmission rate, the distance between communicating devices, the channel characteristics and other factors; there exists no simple relationship. Assuming a fixed distance between transmitter and receiver, a rough estimate of the relative transmit power requirement as a function of the transmission rate to support a desired packet error rate (PER) can be extracted from Fig. 1. For an AWGN channel and a desired PER of 10^{-2}, let the necessary transmit power be $P_{\mathrm{D}}(r_{\mathrm{D}}= 6\text{ Mbps}) = P_0$. The power levels required for the different rates are roughly $P_{\mathrm{D}}(r_{\mathrm{D}}= \{6,9,12,18,24,36,48,54\}\mathrm{Mbps}) = \{1,2,2,4,8,16,32,64\} \cdot P_0$. Hence, a power difference of about 18 dB, or 64 times the power level, is required to maintain the same PER at a rate of 54 Mbps as at 6 Mbps. This power ratio of 18 dB is also confirmed by [24].

The transmit rates r_{A} for the ACK frames are defined by the standard, hence the times T_{A} are given. As the immediate ACK procedure and its reliability are crucial for the system, it is assumed that all ACK frames are transmitted at the highest power level.

From the above results, several observations can be made. If the transmit power is fixed, one should try to transmit at the highest possible rate. This increases the power efficiency and minimizes the air time. This in turn increases the system throughput and minimizes the collision probability. If, on the other hand, the rate is fixed (as might be required by a service), the lowest possible power level should be used to save energy. However, this strategy might worsen the hidden terminal problem. If both the transmit power and data rate are adaptive, the comparison of the transmit time ratio (1/8) and the transmit power ratio (64/1) for 54 Mbps and 6 Mbps suggests to transmit at the lowest possible data rate and then to optimize the transmit power accordingly. However this statement is only valid if we consider the power efficiency for a single communication link. There are several issues that need to be taken into account:

- The power efficiency does not take into account the energy consumption of the baseband and MAC layer signal processing. Higher data rates allow the radio front-ends as well as the PHY and MAC engines to be put into an energy-saving sleep mode for longer time periods.
- The reduced transfer times associated with higher data rates result in a reduced collision probability. This will result in a reduced retransmission rate and better throughput performance than a low data rate transmission.
- Higher power levels associated with higher data rates lead to a less pronounced hidden terminal problem, as more stations can "hear" a station transmitting at a higher power level.
- To optimize the aggregate throughput in a cell, the best strategy is to use the highest possible data rates for data transmission.
- The transfer times depend on the data rate, backoff times, and retransmissions needed. Hence, as long as higher data rates do not considerably increase the retransmission probability, using the highest possible data rates will result in the shortest transfer times.

The arguments given above suggest that the best strategy is to select the highest possible data rate and to adjust the transmit power accordingly.

3 Joint power and data rate adaptation

3.1 Proposed algorithm

The link adaptation algorithm proposed here makes use of the immediate ACK strategy defined in the IEEE 802.11 MAC protocol [3]. Error-free data frames are immediately acknowledged by the receiver. Frames are protected against errors (due to transmission errors or collisions) by means of a frame check sequence (FCS) field containing a 32-bit cyclic redundancy code (CRC) and a simple send-and-wait ARQ mechanism. If the receiver detects a CRC error the frame is discarded. Otherwise, the receiver sends an ACK frame back to the transmitter. If no ACK frame is received within a specified time, the frame is re-sent after a random back-off time. This process is repeated until the transmitter receives an ACK, a packet lifetime value is exceeded, or a maximum number of retries is reached. In the latter two cases, the transmitter discards the frame.

As no channel state information (CSI) is available at the transmitter, the proposed algorithm exploits the immediate ACK procedure's ability to detect frame loss. If the transmitter does not receive an ACK within a specified time interval, the algorithm concludes that link quality was insufficient and that a lower data rate or a higher transmit power should be used. On the other hand, if the transmitter succeeds in sending multiple data frames, it assumes that the link quality is sufficient and that a higher rate or a lower transmit power can be used.

This adaptation scheme can be implemented in the transmitter with a pair of counters for a every destination MAC address: one for successful transmissions s and one for failed transmissions f. If a frame is successfully transmitted, counter s is incremented and counter f reset to zero; similarly, if a transmission fails, counter f is incremented and counter s reset. If the success counter s reaches a certain threshold $S_{\max}$, then the data rate is increased or the power decreased, and both counters are reset to zero. Similarly, if the failure counter f reaches a certain threshold $F_{\max}$, then the data rate is decreased or the power increased, and both counters are reset to zero.

3.2 Dynamic threshold adaptation

The values of parameters $F_{\max}$ and $S_{\max}$ are critical for the performance of the link adaptation scheme [10]. Throughout the paper we assume that $F_{\max} = 1$. This conservative choice corresponds to an immediate adaptation towards a more robust operating point by decreasing the data rate and/or increasing the power in the case of a failed transmission. This strategy prevents unnecessary retransmissions at the expense of a slightly suboptimal steady-state performance.

The optimum choice for $S_{\max}$ depends on the channel dynamics. Fast-changing channels require a small value of $S_{\max}$, so that the transmission parameters can keep up with the channel variations. On the other hand, for slowly changing channels, large values of $S_{\max}$ avoid ineffective switching to higher rates or lower power levels when the channel has not improved.

Exploiting these observations, the basic algorithm described in Section 3.1 is enhanced with a simple but powerful method that estimates the speed of link-quality variations and switches dynamically between two success threshold values, namely S_1 and S_2, with $S_1 < S_2$. In regions of high Doppler spread or fast changing channel conditions, a low threshold ($S_{\max} = S_1$) improves the tracking performance of the algorithm. For low Doppler spread or slowly changing channel conditions a higher threshold ($S_{\max} = S_2$) prevents throughput degradation due to premature switching.

The state transition diagram of the complete adaptation algorithm including dynamic threshold adaptation is shown in Fig. 2. As indicated in the figure, a transmitting node may be in any one of the three states - *High*, *Low*, and *Spread?*. The states *High* and *Low* reflect the current node's assumption about the changing speed of the channel quality. After $S_{\max}$ successful transmissions, the node adjusts its transmission parameters to a higher rate or to a lower power level and enters the state *Spread?*. In this state, it waits for the result of the next transmission to decide whether it should move to the *High* or to the *Low* state. If the next transmission succeeds, the node assumes that the link quality is improving rapidly, enters state *High* and sets $S_{\max}$ to the small value S_1 to react quickly to the changing link quality. If, on the other hand, the next transmission fails, the link quality is assumed to either change slowly or not at all, so that the former decision to change the transmission parameter was premature. Therefore, the state *Low* is entered and $S_{\max}$ is set to the larger value S_2.

3.3 Strategies for data rate and power level selection

From the discussion in Section 2 it is apparent that there is no simple analytical approach for deriving a procedure to select the optimum adaptation parameters, especially in an environment where multiple stations are active. Therefore, we introduce two simple adaptation strategies with focus on (i) compliance with the 802.11 standard and (ii) low implementation complexity. We differentiate between two basic adaptation modes, a “High-Performance” (HP) mode and a “Low-Power” (LP) mode.

In HP mode, the optimization goal is to transmit at the highest possible transmit rate. The transmit power is of secondary importance. This mode is applied when the highest possible rate is required or if the overall throughput should be optimized. The LP mode performs the dual operation to the HP mode. Here, the optimization strategy is to work at the lowest possible power level, and the data rates are adjusted accordingly. It can for example be used in portable devices to save battery power and in cases where the data rate or delay requirements of the application are sufficiently low.

3.3.1 “High-Performance” (HP) mode

The flow diagram for this mode is shown in Fig. 3. As mentioned above, the primary goal in this mode is to support the highest possible data rate. Therefore, after $S_{\max}$ successful transmissions (i.e. it is assumed that the link quality has improved) the data rate is increased until the maximum rate is reached. Once the data rate cannot be increased further, the transmit power level is reduced accordingly. In the case of $F_{\max}$ failed transmissions, the algorithm first attempts to increase the power level to keep the rate at its current value. The rate is only reduced if the transmit power has already reached its maximum level.

A variable $rate_{\text{crit}}$ is introduced to mark a critical data rate, namely, the lowest rate that cannot be supported due to current channel conditions even at the maximum power level. When this variable is set, indicating it does not make sense to increase the data rate to $rate_{\text{crit}}$, it is more suitable to reduce transmit power after $S_{\max}$ successful transmissions. The number of successful power level reductions is tracked in the counter $pcnt$. If the data rate below $rate_{\text{crit}}$ can still be supported after decreasing the power level $P_{\max}$ times, it is assumed that channel conditions have improved. Hence, the algorithm increases the data rate to $rate_{\text{crit}}$ using the highest possible transmit power level and $rate_{\text{crit}}$ is incremented or reset to zero to allow operation at the maximum data rate.

3.3.2 “Low Power” (LP) mode

In LP mode, we try to use the lowest possible transmit power level, and only adjust the transmit rate when we can no longer decrease the power level. The flow diagram for this mode can be derived from Fig. 3, as it performs the dual operation of the HP mode with the following dualities: *rate* is replaced by *pow*, the increment operator $++$ by the decrement operator $--$, and min by max. In addition, instead of locking the rate $rate_{\text{crit}}$ the power pow_{crit} is locked, *pcnt* is replaced by the rate counter *rcnt*, and the counter limit $P_{\max}$ is replaced by $R_{\max}$.

The key operations in LP mode can be described as follows: when a certain power level does not result in a successful transmission even for the lowest possible data rate, then the unsuccessful power level is stored in pow_{crit} and the power level is increased. If the power level above the critical, currently unsupported, level pow_{crit} still results in successful transmissions after increasing the data rate $R_{\max}$ times, it is assumed that the channel conditions have improved. Hence, the algorithm tries to decrease the power level to pow_{crit} using the lowest possible data rate. To further decrease the minimum power level, the critical value pow_{crit} is either decreased or reset to zero to allow operation at the minimum power level.

The two strategies described above for HP and LP mode, respectively, rely on the assumption that the discrete step sizes available for selecting the data rate and adjusting the power level provide sufficient optimization potential to move the operating point as closely as possible to the optimum power efficiency value. This will be verified in Section 4 by means of simulation results.

3.3.3 Single-parameter link adaptation

Two special cases can be obtained if either the data rate or the power level is fixed. In [10] a "Rate-Only" (RO) mode was studied, in which the transmit power level is fixed and the data rate is adjusted using the dynamic threshold adaptation scheme shown in Fig. 2. Another possible optimization strategy was considered in [22]. Here, the data rate is fixed and the power level is adjusted. The data rate is increased only if the lowest possible power level has been reached. The performance of single-parameter link adaptation will be addressed in Section 4, together with that of the proposed joint data rate and power optimization schemes.

4 Performance Evaluation

To study and understand the behavior and performance of WLANs with the link adaptation algorithm proposed in this paper, we have implemented a discrete-event simulator which models the PHY and MAC layers of 802.11a WLANs in detail. A comprehensive description of the simulator is given in [10].

We focus first on a simple WLAN link with only one AP and one WSTA (point-to-point WLAN), and study the impact of the various parameters of the link adaptation algorithm on its throughput and delay performance. Then, infrastructure WLANs with multiple WSTAs will be considered, in Section 4.2 for Poisson traffic sources, and in Section 4.3 with more realistic source models for data, telephony, and video services. Furthermore we distinguish between saturated and non-saturated networks. In a saturated network the traffic at the nodes is generated in such a way that the nodes always have packets to send. In such a network the throughput is the most important performance value. On the other hand, in a non-saturated network, the total generated load is rather small compared to the capacity of the WLAN, and delays experienced by the data packets are more important than the achievable throughput.

Unless stated otherwise, the following parameter values are used for the simulation runs:

- Fixed Doppler spread of $f_d = 5$ Hz (which is consistent with measurement results in office buildings [25]);
- constant packet length of 1000 bytes;
- negative-exponential distributed packet arrivals;
- transmit power levels max = $+10$ dBm, min = -10 dBm;
- transmit power step sizes up = $+5$ dB, down = -2 dB (the choice of these parameter values will be explained later on);
- success thresholds $S_1 = 3$ and $S_2 = 10$; and
- failure threshold $F_{max} = 1$.

4.1 Point-to-point wireless LAN

4.1.1 Saturated point-to-point wireless LAN

In this section we study the performance of an 802.11a WLAN with two stations as function of the distance between the two nodes and the adaptation schemes. The diagrams are obtained by having both stations saturated with a traffic load of 30 Mbps. For all adaptation schemes the transmit power level is fixed at $+10$ dBm, except for the HP and LP algorithms, for which it varies between $\pm$ 10 dBm with step widths of $+5$ and -2 dB. As reference the curves for the RO and the “genie-aided” adaptation schemes are also shown in the diagrams. The RO scheme is the one in which only the data rates are adapted, while the power level P_D remains fixed. The performance of this mode was studied in detail in [10]. In particular it was shown in [10] that this mode provides performance values very close to those obtained with a “genie-aided” adaptation scheme in which the transmitter knows the CSI a priori.

The throughput of this point-to-point WLAN is shown in Fig. 4. As expected, it decreases with increasing distance between the nodes; furthermore, low data rates support larger ranges than high data rates do. It can be seen that the HP algorithm (upper solid line) provides throughput values very close to those of the “genie-aided” and the RO modes (dashed lines). The small throughput degradation when compared to those two modes is due to its permanent attempt to reduce the transmit power level. This is also the reason it provides less throughput than the 54 Mbps fixed-rate mode in the case of a short distance. Otherwise, it always achieves much higher throughput than any of the fixed-rate modes do.

In addition, we note that the LP mode (lower solid line) achieves high throughput values only at short distances. Its throughput performance degrades rapidly with increasing distances. However, it remains as good as the 6 Mbps fixed-rate mode at large distances. This behavior results from its strategy of using the lowest possible transmit power level and only optimize (i.e. increase) the transmit rate when the minimum power level has been reached.

It is interesting to observe how the throughput efficiency of the fixed-rate modes deteriorates with higher data rates, even when the distances are short. Whereas the 6 Mbps fixed-rate mode achieves a maximum throughput of approximately 5 Mbps, which corresponds to an efficiency of approximately 83%, the efficiency of the 24 Mbps rate is reduced to approximately 70%, and that of the 54 Mbps drops to 50%. This is caused mainly by the backoff and ACK procedures of the 802.11 MAC layer: these overhead procedures are independent of the rates used. At high data rates the transmission time of a MAC frame is reduced but not those of the overhead procedures, thus leading to a higher throughput loss. An exhaustive discussion of the impact of the MAC overhead on the 802.11 throughput performance can be found in [23].

Fig. 5 demonstrates how well the adaptation algorithms can manage the transmit power problem. In particular, at distances smaller than 20 m, both the HP and LP modes (solid lines) reduce the transmit energy significantly compared to the RO (dashed line) and the fixed-rate modes (dotted lines). As mentioned in Section 2 the power consumed during the idles times (e.g. $DIFS$, $SIFS$) is set to zero in all simulation results. Thus the average transmit power $\bar{P}_{tx}$ displayed in Fig. 5 contains only the energy consumed during frame transmissions. For the RO mode, $\bar{P}_{tx}$ first increases with increasing distance then remains almost constant. As the transmit power for that mode is fixed to +10 dBm, the initial increase of $\bar{P}_{tx}$ is due to the switch to lower transmit rates at higher distances, which results in longer transmission times and thus in higher energy consumption. Similarly, the increase of $\bar{P}_{tx}$ for the HP and LP modes (solid lines) is caused by both rate reduction and power increase to support larger distances. As expected, at long distances, the $\bar{P}_{tx}$ of the HP, LP, and RO modes approach the values of the 6 Mbps fixed-rate mode because the adapted transmission rate has reached the lowest value of 6 Mbps and the transmit power level the highest value of +10 dBm. There is simply no adaptation potential left.

It is interesting to observe the general behavior of the fixed-rate modes. With increasing distance their $\bar{P}_{tx}$ values first decrease and then settle at constant values at distances where their throughput approaches zero. This behavior corresponds to the fact that with increasing distance between nodes the probability of erroneous transmissions increases. Consequently, nodes spend more time in the backoff mode than in the transmit mode; hence the average transmit power is reduced. At very large distances, almost all transmissions are erroneous, leading to a constant average transmit power in the regions where the throughput is zero.

An attentive reader may notice the slight decrease of the average transmit power for the LP mode at distances around 10 m. When the distance between the two stations is very short, a high throughput is achieved by means of the highest data rate combined with the lowest transmit power level for the data frames. However, as mentioned before, we select the strategy of sending the ACK frames with the highest possible power level because of their vital importance for the adaptation algorithms; we want to avoid retransmissions of data frames that are caused by erroneous ACK frames. At very short distances, the energy used to send ACK frames outweighs the one used to send data frames. The ratio of this ACK energy fraction to the one used to transmit data frame is then reduced as the distance increases and lower data rates (i.e. larger transmission times and higher energy consumption) have to be used.

The power efficiencies η_{TP} of the various adaptation modes are compared in Fig. 6. The higher the η_{TP} is, the more efficient is the corresponding adaptation mode. As expected, the LP mode achieves the best power efficiency, followed by the HP mode. Both schemes are very efficient, in particular for distances up to 20 m. In this region the LP mode outperforms the HP mode because of its strategy of first reducing the power level and only increasing the rate once the minimum power level has been reached. This superiority however comes at the price of very fast throughput degradation with distance, as shown in Fig. 4. When the two nodes are very close to each other, the efficiency of both modes is almost the same, because in this case both modes operate at the maximum transfer rate and the lowest power level. Compared with the HP and LP modes, the remaining schemes are less efficient.

Until now we have fixed the power step sizes used by the HP and LP modes, namely to $+5$ dB when increasing the power level (the *up* size) and to -2 dB when decreasing it (the *down* size). Fig. 7 shows the impact of different step sizes on the performance of the wireless link. Whereas the throughput achieved by the LP mode changes only marginally when the step sizes are varied, the performance of the HP mode could be negatively affected by an improper step size selection. Taking the curve with the equal up-down sizes of ± 2 dB as reference (solid line), we can see that a high down size of -5 dB dramatically reduces the throughput of the link (dotted line), whereas the same value in the opposite direction (up = $+5$ dB, down = -2 dB) improves the performance (dashed line). This behavior is due to the fact that a large decrease in transmit power level results in additional transmission errors, whereas a large power level improves the probability of successful transmissions. A further increase of the up size to $+8$ dB does not bring further improvement, because the power levels used in the system considered were limited to ± 10 dBm.

4.1.2 Non-saturated point-to-point wireless LAN

In this section we study the performance of the various adaptation modes if the point-to-point wireless link is not saturated. The distance between the two nodes is fixed to 20 m, and the total offered load generated at both stations varies between 0 and 20 Mbps. We are interested in the transfer times of the data packets, which is an important performance measurement for delay-sensitive applications such as telephony or interactive video. The transfer time is defined as the duration from the packet's arrival at the transmitting node until its successful receipt by the destination node. Thus, it includes the waiting time the packet experiences at the source node and all its transmission and re-transmission times.

The results for the average transfer time are shown in Fig. 8. It can be observed that the RO mode (dashed line) has the best delay performance followed by the HP mode. Whereas the fixed-rate modes can at best support an offered load of 16 Mbps before becoming overloaded (see the curve for the 36

Mbps fixed-rate mode), the HP and RO adaptation schemes still provide finite mean transfer times for loads as high as 20 Mbps. The LP mode can only supports loads up to approximately 4 Mbps. It is interesting that for the channel type considered here the 24 Mbps fixed-rate mode achieves lower delay values than the 36 or 54 Mbps ones, even for rather low loads from 2 to 14 Mbps. This behavior is caused by the large variance of the transmission time at high data rates. As the data rate increases, the transmission time of a frame becomes shorter if the frame is sent without error. However, if the frame is disturbed, its transmission time is increased by backoff times, which are independent of the data rates. Hence the variance of the transmission time increases with higher data rates, which in turn leads to larger mean waiting and mean transfer times.

The average transmit power and the power efficiency of the point-to-point wireless link are shown in Fig. 9. Again, similar to the non-saturated case, the LP mode provides the best power efficiency at the price of high transfer delays. The HP and RO schemes achieve almost the same efficiency at low offered loads. The advantage of the HP mode becomes more apparent in the region of higher loads; the more data are to be sent, the better the HP mode can adapt the transmission parameters to the changing link states.

4.2 Infrastructure wireless LAN with multiple WSTAs

Until now we have studied the performance of a rather simple point-to-point WLAN. In the following subsections we increase the system complexity by considering WLANs with multiple WSTAs.

4.2.1 Saturated wireless LAN

Fig. 10 shows the throughput performance of a saturated 802.11a WLAN with 10 WSTAs placed randomly within a circle of 60 m radius and the AP in the center. A traffic load of 3 Mbps is generated at each WSTA, and no traffic is generated at the AP. For each adaptation mode, 20 independent simulation runs were performed using different random generator seeds. All WSTAs employ the same adaptation scheme during a given simulation run. Each point in Fig. 10 represents one WSTA. There are 20 runs $*$10 WSTAs = 200 points for each adaptation mode.

It can be observed that for the 54 Mbps fixed-rate mode as one extreme case, the generated load of 3 Mbps can only be supported by stations that are not more than 30 m away from the AP. WSTAs located farther away cannot communicate with the AP. In the other extreme case, where all WSTAs communicate with the 6 Mbps fixed-rate mode, all stations can still communicate with the AP but they achieve only a throughput of approximately 0.5 Mbps. The LP mode shows a slightly better throughput than the 6 Mbps fixed-rate mode, the HP improves the throughput further, and the RO mode has the best performance. For distances smaller than 30 m, the three modes HP, LP, and RO achieve lower throughput values than the 54 Mbps fixed-rate one. This throughput reduction is because WSTAs that are located far away from the AP can now communicate when using these adaptation schemes. Thus, the total throughput of the system is more fairly distributed to all WSTAs by the three modes HP, LP, and RO than by the fixed-rate modes.

Compared with the power efficiency results obtained for the point-to-point case shown in Fig. 6, Fig. 10(b) illustrates that while all other schemes achieve almost the same power efficiency as in the point-to-point scenario, the LP mode suffers a significant efficiency degradation for stations that are close to the AP. This performance loss is due to the well-known hidden station issue. Stations that are close to the AP transmit with low power. Therefore their transmissions cannot be detected by the other stations, which results in an increase in the number of collisions.

4.2.2 Non-saturated wireless LAN

The transfer delay performance of a non-saturated 802.11a infrastructure WLAN with 10 WSTAs distributed within a circle of 20 m is shown in Fig. 11. Each WSTA carries a traffic load of 0.3 Mbps. Hence, even the 6 Mbps fixed-rate mode should be able to support the total generated load of 3 Mbps. As can be seen from Fig. 11(a), the shortest transfer delay is provided by the RO scheme, followed by HP and the 24 Mbps fixed-rate modes. The results for the 54 Mbps mode shows a steep increase of the mean transfer time with increasing distance from the AP. The LP mode provides low transfer delays only to WSTAs that are close to AP; otherwise it achieves almost as long delays as the 6 Mbps fixed-rate mode, but with a significantly lower average transmit power, see Fig. 11(b). It is interesting that the hidden station issue as experienced in the former scenario (Fig. 10(b)) is not apparent in this figure. The reasons are that the WSTAs now carry only $1/10$ of the traffic load, i.e. the probability that two or more WSTAs have data to send at the same time is smaller.

4.3 Infrastructure wireless LAN with realistic traffic models

So far all simulation runs were done with traffic generated by a Poisson process for the packet arrivals. In this section we will consider a scenario with more realistic traffic sources. It still is an infrastructure WLAN with 1 AP and 10 WSTAs distributed within a circle of 20 m radius, but now the stations support three types of services, namely best-effort data, interactive telephony, and real-time video. The best-effort data service is supported by all WSTAs and is modeled by sources generating packets having exponentially distributed inter-arrival times. The packet-length distribution is taken from a real LAN trace, with a mean packet length of 501 bytes [26]. An up/downlink ratio of 20 Kbps/100 Kbps has been chosen to emulate the asymmetrical behavior of Web-browsing-like services.

Besides the data service, one half of the WSTAs supports telephony and the other half video services. Telephony voice signals are modeled by a two-state ON/OFF process, emulating talk spurts followed by silence periods. The talk-spurt and silence duration times are exponentially distributed with a mean value of 1 s and 1.35 s, respectively [27]. During the talk spurts, the voice signals are encoded with the ITU-T G.711 speech codec [28], which generates 160-byte-long packets every 20 ms, corresponding to a constant bit-rate of 64 kbps. Furthermore, a header compression has been assumed that reduces the 40 bytes of the RTP/UDP/IP header to 4 bytes [29].

The video service emulates a MPEG downlink streaming service delivered to the end user from a storage database, e.g. a DVD player or a Web Video Content Server. We assume that the video flows are generated by a CBR (Constant Bit Rate) source with a rate of 500 kbps and a fixed packet length of 1500 bytes.

Fig. 12 displays the uplink and downlink mean transfer times for the voice service. It can be seen that the RO mode achieves the best delay performance, followed by the HP and the 24 Mbps fixed-rate schemes. The LP mode provides similar results as the 6 Mbps fixed-rate mode. While the 54 Mbps fixed-rate scheme exhibits quite small values in the uplink direction, its results for the downlink direction are not visible in Fig. 12(b) because of their very high values. In general, the downlink performs worse than the uplink because of the nature of the infrastructure WLAN in which all traffic has to go via the AP, thus the AP has to carry a higher traffic load than the WSTAs. Note that the points on the x-axis of Fig. 12 correspond to those WSTAs that do not have telephony service. The same applies to Fig. 14.

The results for the throughput of the best-effort data service are shown in Figs. 13(a) and 13(b). With the exception of 54 Mbps fixed-rate scheme, all modes are capable of supporting the offered data traffic for distances up to 20 m. Similar results are obtained for the throughput of the downlink video stream as shown in Fig. 14, in which the video throughput is slightly reduced when the WLAN operates in the 54 Mbps fixed-rate mode.

The average transmit power and the power efficiency of the WSTAs are shown in Fig. 15. Again, it can be seen that in terms of power efficiency the two adaptation modes HP and LP provide the best values compared to the other modes without power regulation. Note that for each mode there are two sets of points, corresponding to WSTAs with and without telephony traffic; here, WSTAs with telephony traffic have a better efficiency. The more data a station has to send, the more efficiently it works. Furthermore, those that receive video streams spend a lot of power acknowledging the video frames coming from the AP.

5 Conclusion

In this paper, we have introduced new, simple but powerful dynamic link adaptation schemes for 802.11a and 802.11g WLANs. These schemes utilize the immediate ACK strategy for adjusting the link parameters. Hence, the algorithms work autonomously in the transmitter; no additional feedback is required from the receiver side, which in turn provides a fully standard-compliant solution. The required additional implementation complexity is very low: Because the approach only counts successful and unsuccessful transmissions, there is no channel estimation or other expensive operation involved. Compared with most of the approaches discussed at the beginning of the paper, no RTS/CTS scheme is required here. It can however be applied as an independent option if necessary.

Throughout the paper, we have analyzed several possibilities to adapt the link parameters power and data rate dynamically in dependence of the channel conditions. It has been shown that the optimization strategy depends on the scenario considered, the number of active stations in a system and the service requirements. However, it turns out that in many cases the strategy to transmit at the highest possible rates with rather high power levels, as is done in the HP mode, is the best choice. The most important arguments to support this statement are the achievable low transfer times, the supported higher overall throughput, and the relaxation of the hidden terminal problem. Moreover, when also taking into account the overall power consumption during idle times, the strategy to keep the transmission times as small as possible provides the greatest potential for putting idle components into a low-power mode.

The HP mode suffers some throughput degradation compared with the RO mode because of unsuccessful switching attempts; it is however much more power-efficient for low distances between the transmitter and the receiver. The LP mode should be considered only if energy is the most critical issue and the required data rates are low.

References

[1] IEEE Std 802.11a-1999. Supplement to IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High-speed Physical Layer in the 5 GHz Band, 1999.

[2] IEEE Std 802.11g-2003. IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 4: Further Higher Data Rate Extension in the 2.4 GHz Band, June 2003.

[3] IEEE Std 802.11-1999. Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 1999.

[4] G. Holland, N. Vaidya, and P. Bahl. A rate-adaptive MAC protocol for multi-hop wireless networks. In *Proc. MOBICOM'01*, Rome, Italy, July 2001.

[5] A. Kamerman and L. Montean. WaveLAN-II: A high-performance wireless LAN for the unlicensed band. *Bell Labs Technical Journal*, pages 118–133, Summer 1997.

[6] D. Qiao and S. Choi. Goodput enhancement of IEEE 802.11a wireless LAN via link adaptation. In *Proc. ICC'2001*, Helsinki, Finnland, 2001.

[7] A. Forenza and R.W. Heath Jr. Link adaptation and channel prediction in wireless OFDM systems. In *Proc. Globecom'02*, 2002.

[8] S.T. Sheu, Y.H. Lee, and M.H. Chen. Providing multiple data rates in infrastructure wireless networks. In *Proc. Global Telecommunications Conf. 2001*, volume 3, pages 1908–1912, San Antonio, TX, USA, 2001.

[9] Y.J. Cheng, Y.H. Lee, and S.T. Sheu. Multi-rate transmissions in infrastructure wireless LAN based on IEEE 802.11b protocol. In *Proc. 54th VTC 2001*, volume 4, pages 2609–2612, Atlantic City, NJ, USA, 2001.

[10] A. Noll Barreto, P. Chevillat, J. Jelitto, and H.L. Truong. A dynamic link adaptation algorithm for IEEE 802.11a wireless LANs. In *Proc. ICC 2003 - IEEE 2003 Int'l Conf. on Communications*, volume 2, pages 1141–1145, Anchorage, AK, USA, May 2003.

[11] J. del Prado Pavon and S. Choi. Link adaptation strategy for IEEE 802.11 WLAN via received signal stregth measurement. In *Proc. ICC 2003 - IEEE 2003 Int'l Conf. on Communications*, Anchorage, AK, USA, May 2003.

[12] M. Krunz and A. Muqattash. A power control scheme for MANETs with improved throughput and energy consumption. In *Proc. of the 5th International Symposium on Wireless Personal Multimedia Communications*, volume 2, October 2002.

[13] C.R. Lin and C.Y. Liu. Enhancing the performance of IEEE 802.11 wireless LAN by using a distributed cycle stealing mechannism. In *Proc. 4th IEEE Conference on Mobile and Wireless Communications Network*, Stockholm, Sweden, September 2002.

[14] S. Agarwal et al. Distributed power control in ad hoc wireless networks. In *Proc. IEEE PIMRC'01*, 2001.

[15] J.-P. Ebert and A. Wolisz. Combined tuning of RF power and medium access control for WLANs. *Mobile Networks & Applications*, 5(6):417–426, September 2001.

[16] A. Kraemling et al. Interaction of power control and link adaptation for capacity enhancement and QoS assistance. In *Proc. of 13th IEEE International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC 2002)*, volume 2, Lisbon, Portugal, September 2002.

[17] Daji Qiao, Sunghyun Choi, Amjad Soomro, and Kang G. Shin. Energy-efficient PCF operation of IEEE 802.11a wireless LAN. In *Proc. IEEE INFOCOM02*, New York, USA, June 2002.

[18] D. Qiao, S. Choi, A. Jain, and K.G. Shin. Adaptive transmit power control in IEEE 802.11a wireless LANs. In *Proc. 57th IEEE Semianual Vehicular Technology Conference*, Jeju, Korea, April 2003.

[19] S. Simoens and D. Bartolom. Optimum performance of link adaptation in HIPERLAN/2 networks. In *Proc. of VTC 2001*, 2001.

[20] Z. Lin, G. Malmgren, and J. Torsner. System performance analysis of link adaptation in HIPERLAN type 2. In *Proc. of the VTC 2000 Fall*, 2000.

[21] S.L. Wu, Y.C. Tseng, C.Y. Lin, and J.P. Sheu. A multi-channel MAC protocol with power control for multi-hop mobile ad hoc networks. *The Computer J.*, 45(1):101–110, 2002.

[22] J. Jelitto, A. Noll Barreto, and H.L. Truong. Power and rate adaptation in IEEE 802.11a wireless LANs. In *Proc. 57th IEEE Semianual Vehicular Technology Conference*, Jeju, Korea, April 2003.

[23] J. Jelitto and H. Linh Truong. A multi-MAC approach to overcome the throughput limits of 802.11 WLANs. In *Proc. 2004 International Zurich Seminar on Communications*, Zurich, Switzerland, February 2004.

[24] R. van Nee and R. Prasad. *OFDM Wireless Multimedia Communications*, chapter 10.5.6. Artech House, Boston, London, 2000.

[25] H. Hashemi. The indoor radio propagation channel. *Proc. IEEE*, pages 943–968, July 1993.

[26] B.H. Walke et al. IP over wireless mobile ATM - guaranted wireless QoS by HiperLAN/2. *Proc. IEEE*, 89:21–40, January 2001.

[27] D.J. Goodman and H. Xie. PRMA system performance with variable number of calls. *Technical Report WINLAB-TR-36*, June 1992.

[28] ITU-T Recommendation G.711. Pulse Code Modulation (PCM) of Voice Frequencies, 1988.

[29] K. Svanbro. Lower layer guidelines for robust RTP/UDP/IP header compression. *RFC 3409*, December 2002.

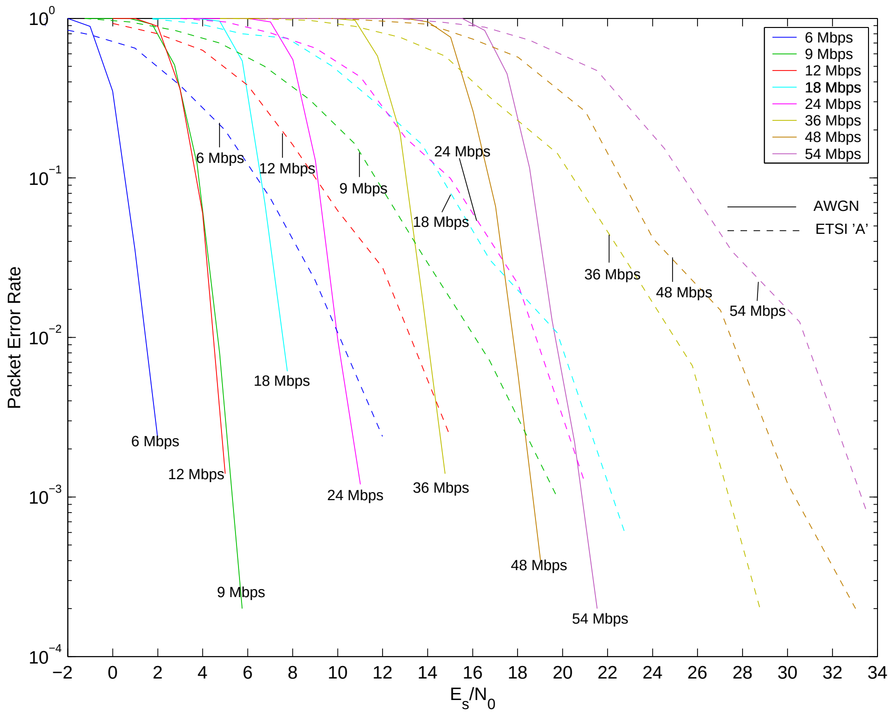

Figure 1: Packet error rate as function of E_s/N_0 and the 802.11a data rates for AWGN (solid) and ETSI ‘A’ (dashed) channel (18 taps, Rayleigh fading, $\sigma_{rms} = 50$ ns, $\sigma_{max} = 400$ ns).

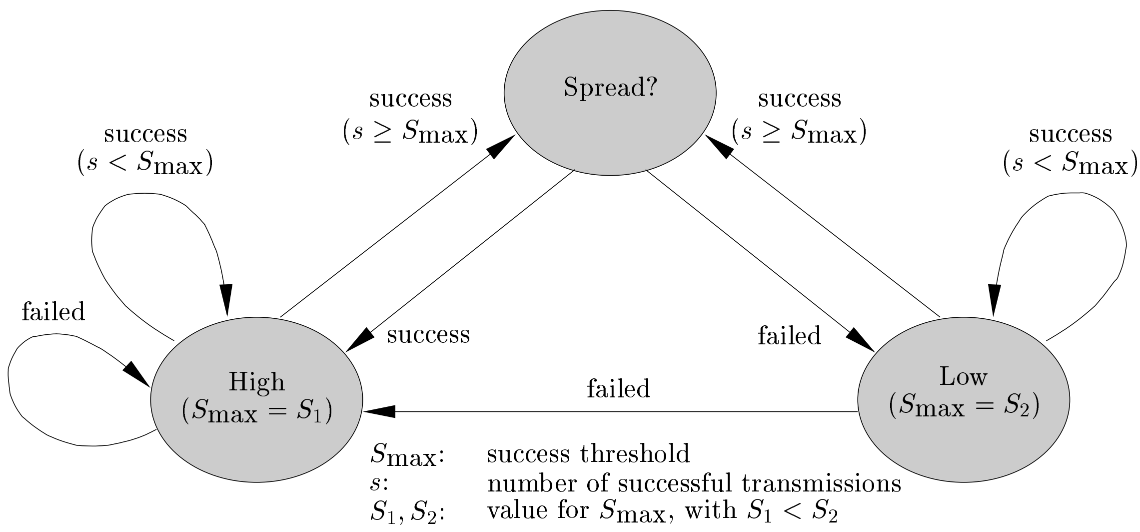

Figure 2: Transition diagram of the link adaptation algorithm.

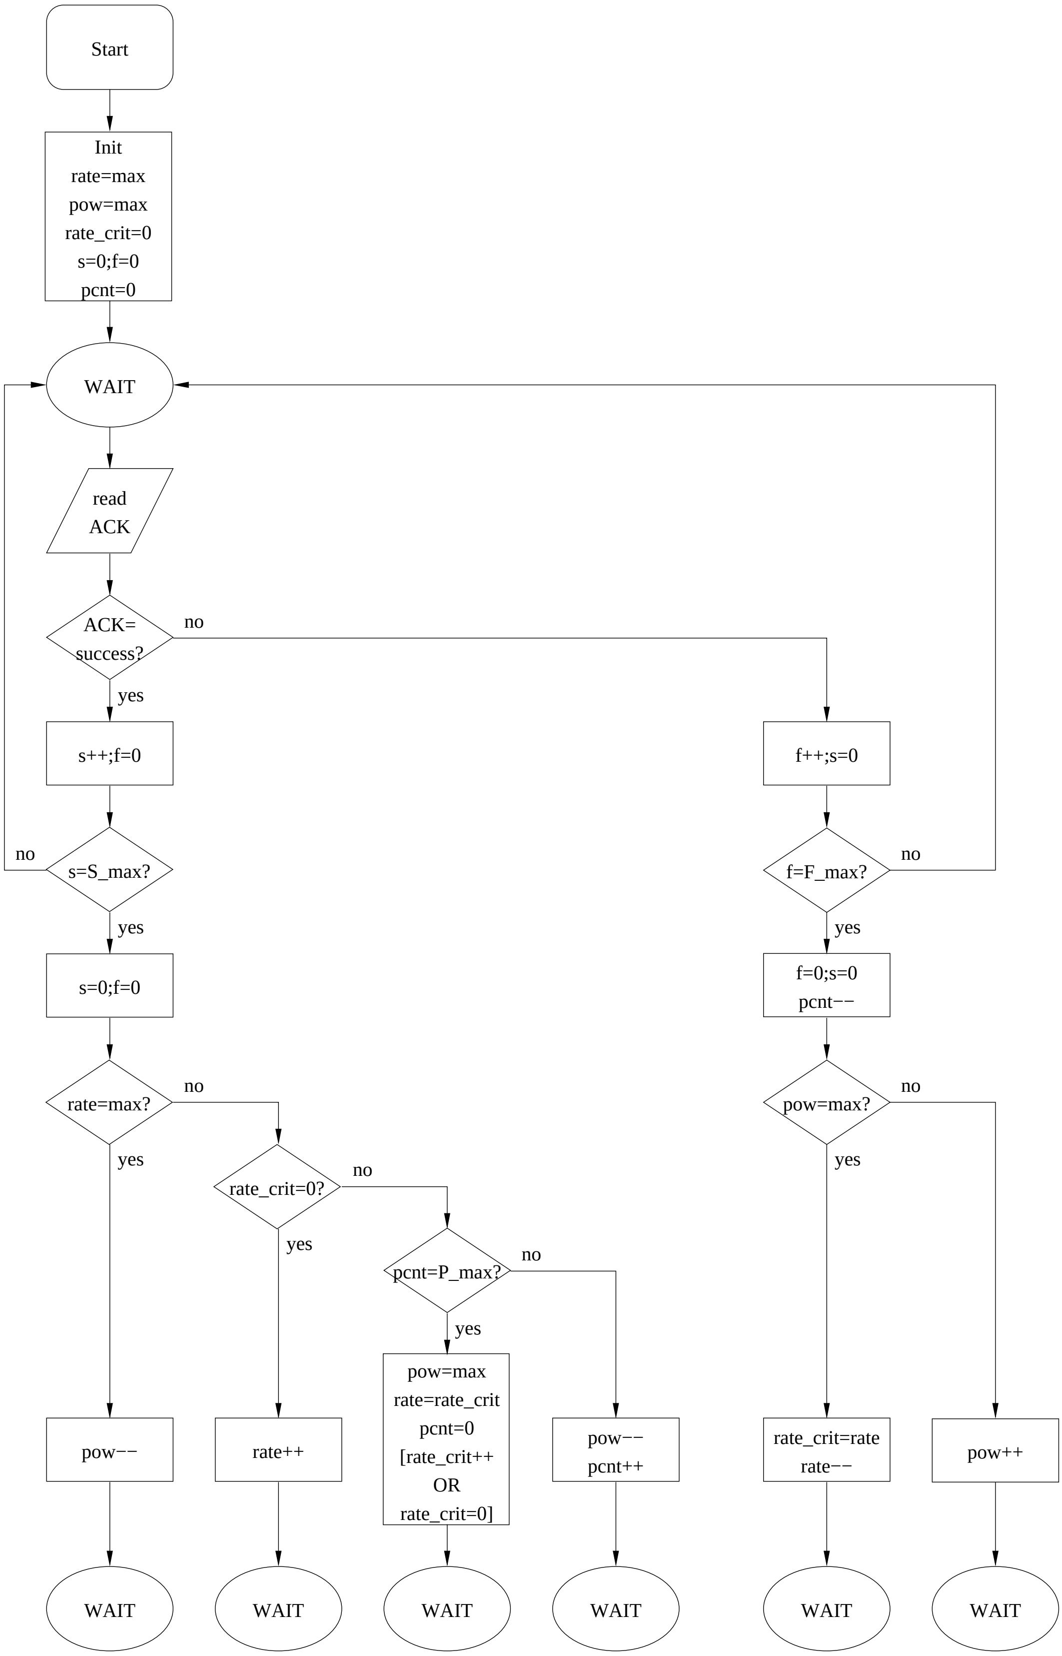

Figure 3: Joint power and rate adaptation algorithm. Flow diagram for HP mode.

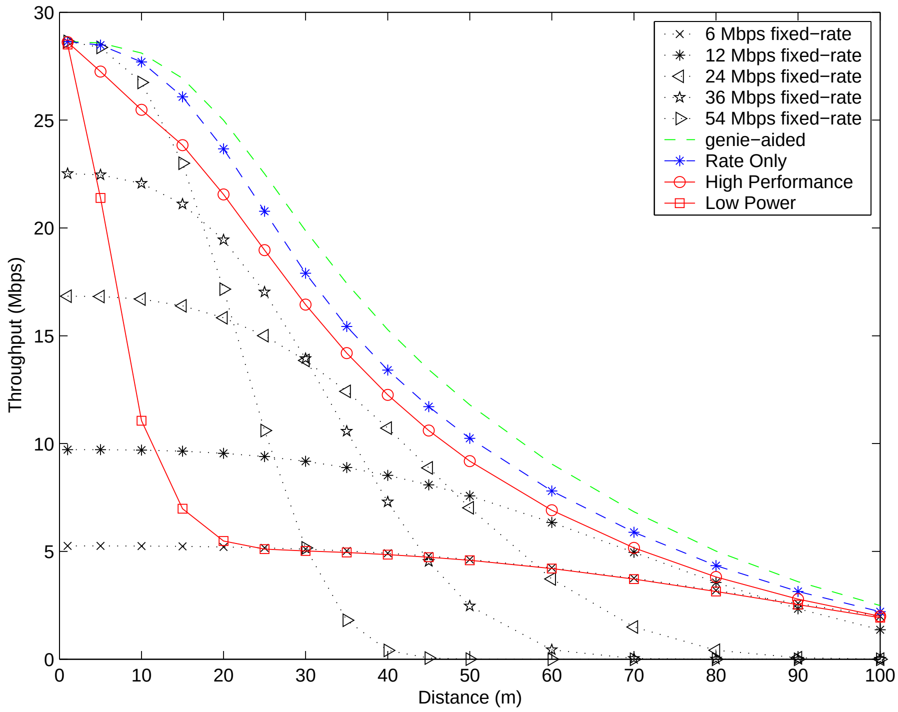

Figure 4: Throughput in a saturated 802.11a point-to-point WLAN as function of the distance between the two nodes.

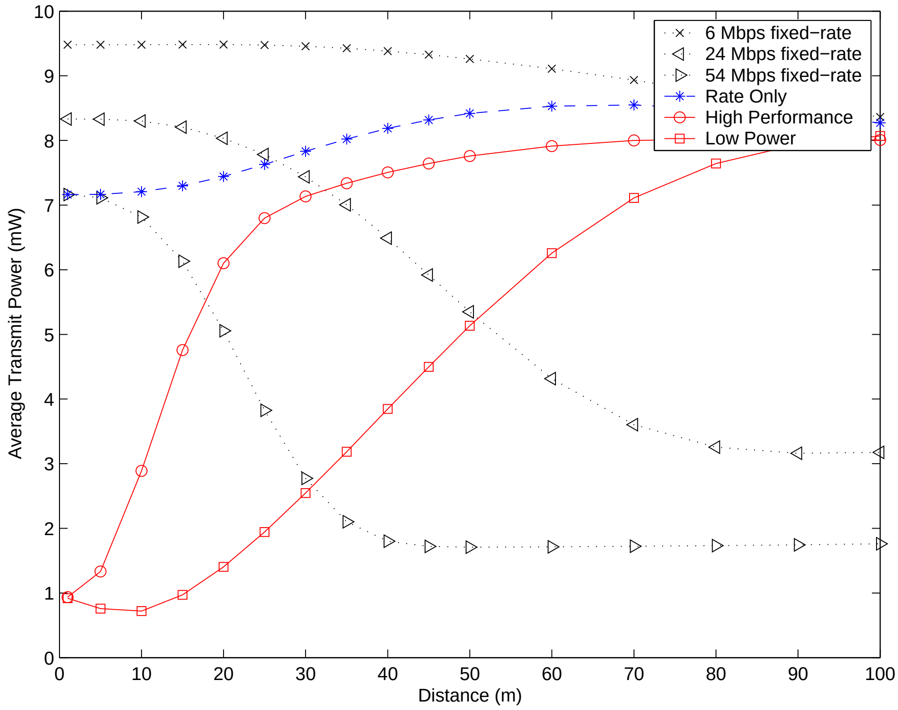

Figure 5: Average transmit power in a saturated 802.11a point-to-point WLAN as function of the distance between the two nodes.

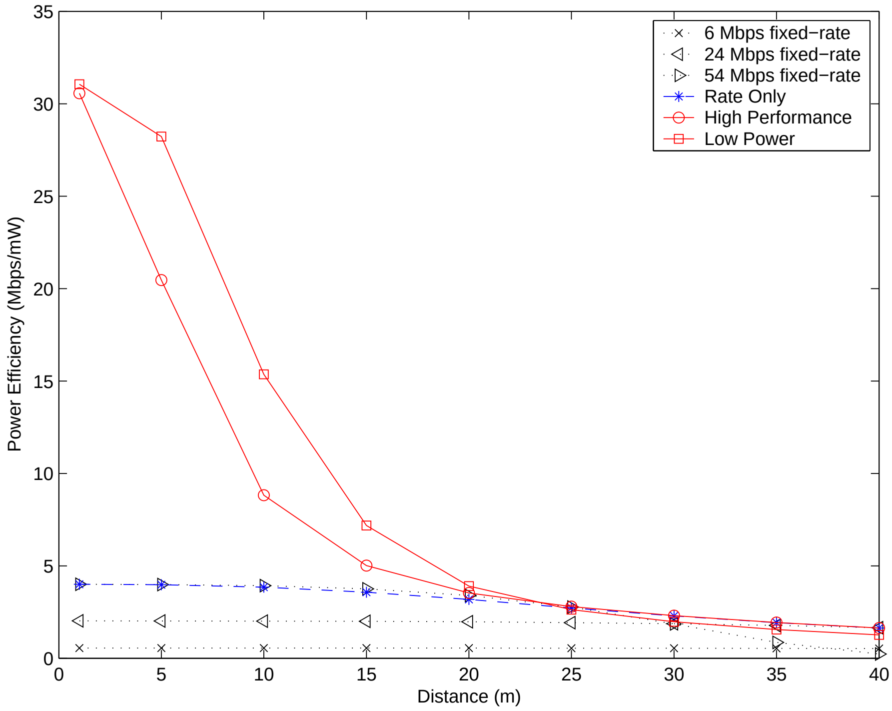

Figure 6: Power efficiency in a saturated 802.11a point-to-point WLAN as function of the distance between the two nodes.

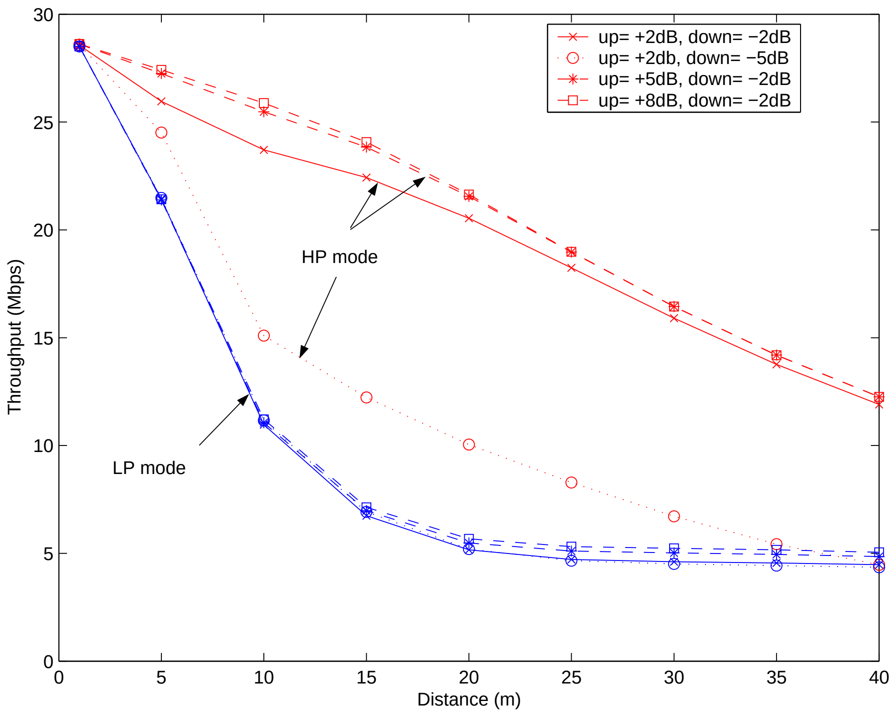

Figure 7: Throughput of a saturated 802.11a point-to-point WLAN as a function of the distance between the two nodes. Impact of the power step sizes.

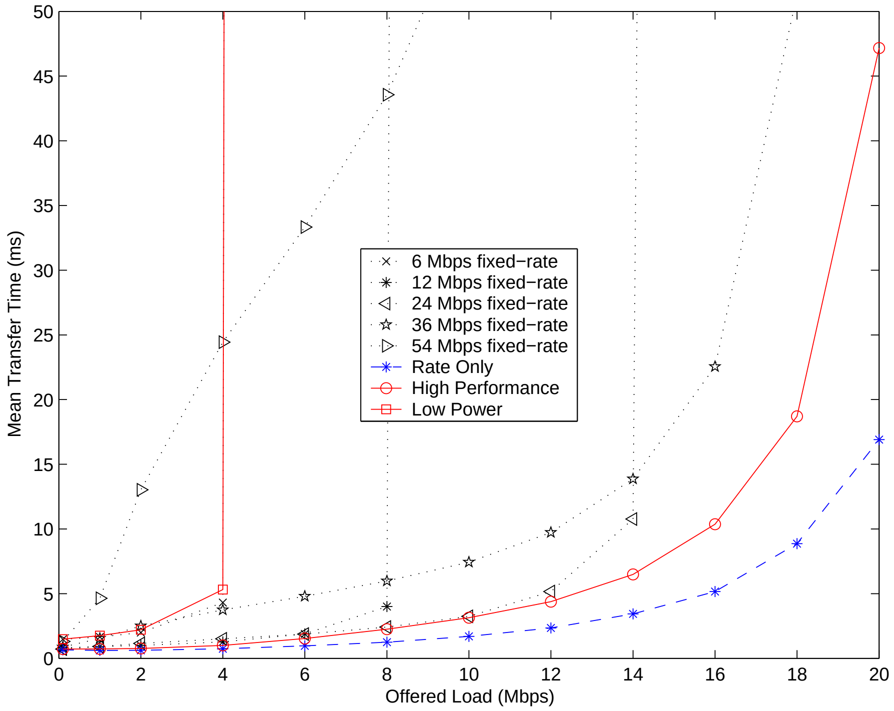

Figure 8: Mean transfer time in an 802.11a point-to-point WLAN as function of the offered load (20 m distance between the two nodes).

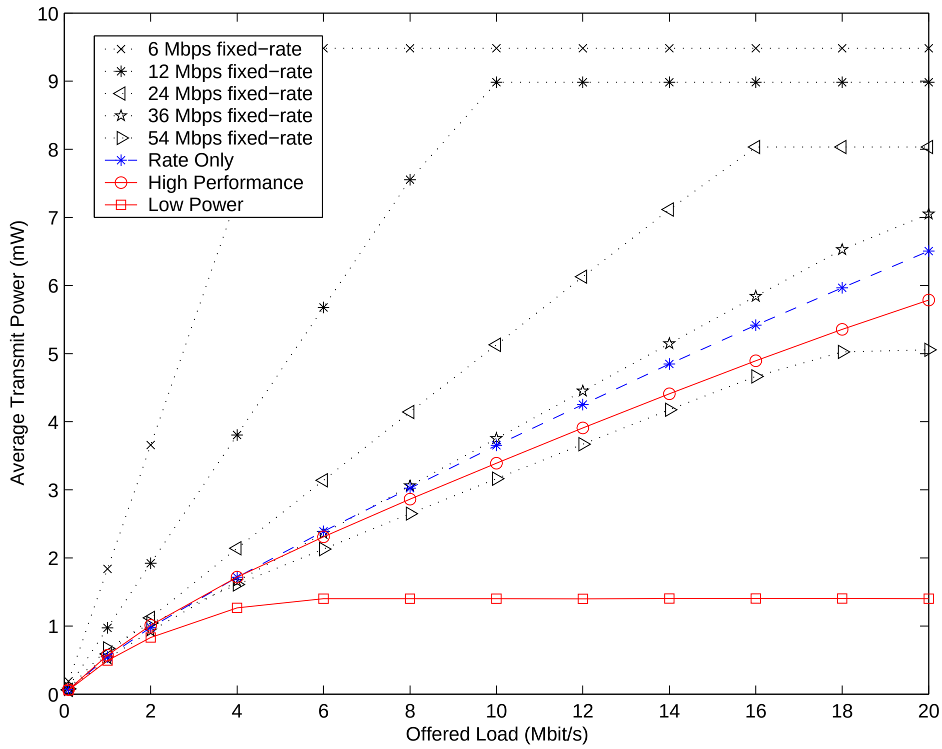

(a)

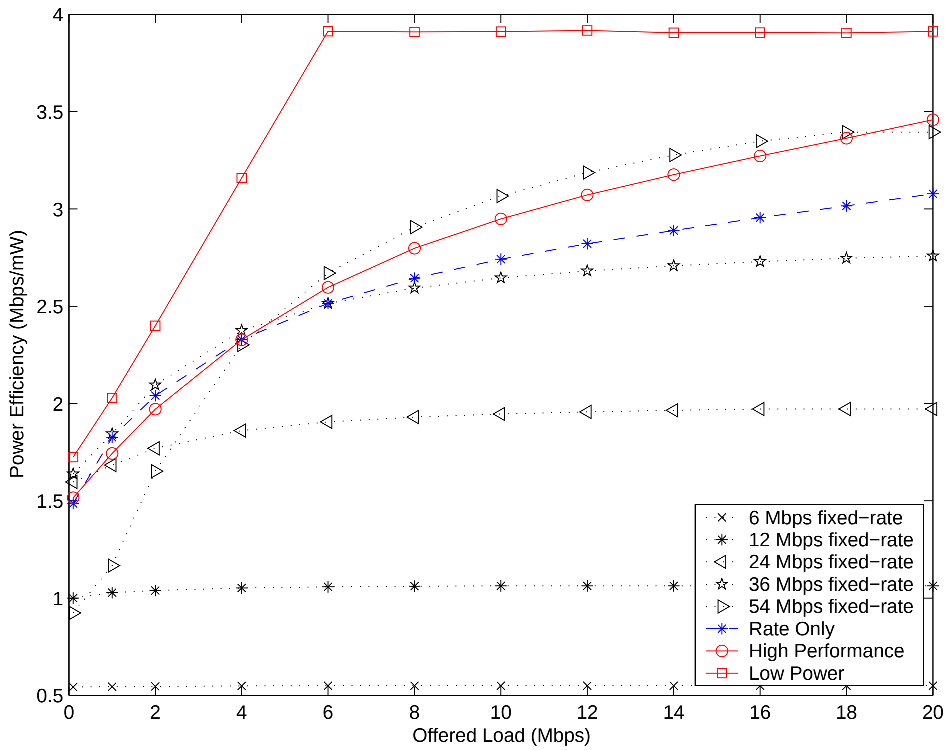

(b)

Figure 9: (a) Average transmit power and (b) power efficiency in an 802.11a point-to-point WLAN as function of the offered load (20 m distance between the two nodes).

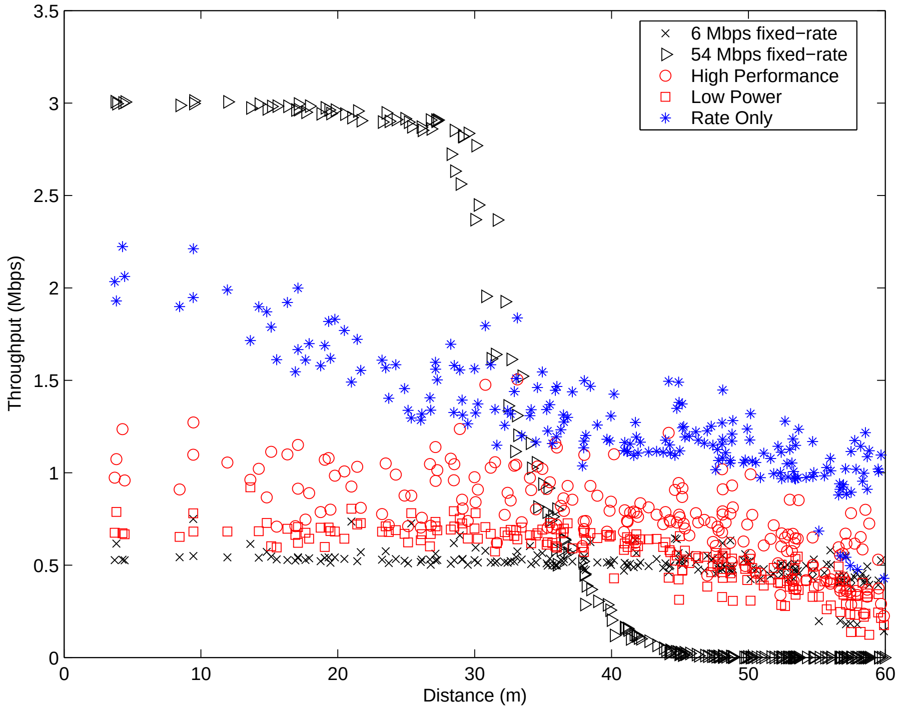

(a)

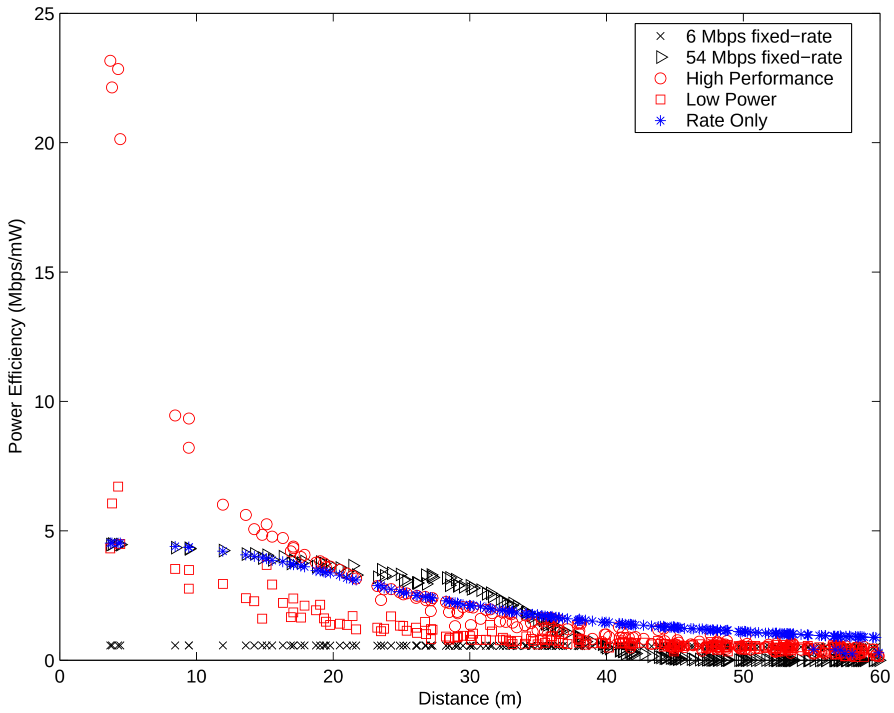

(b)

Figure 10: (a) Throughput and (b) power efficiency in an 802.11a infrastructure WLAN with 10 WSTAs as function of the distance to the AP.

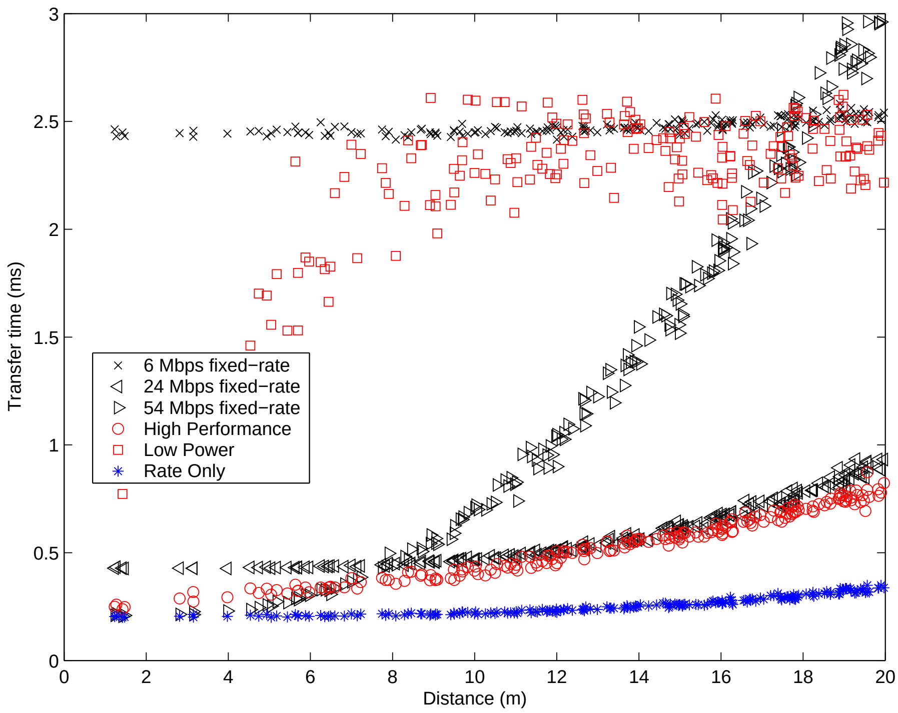

(a)

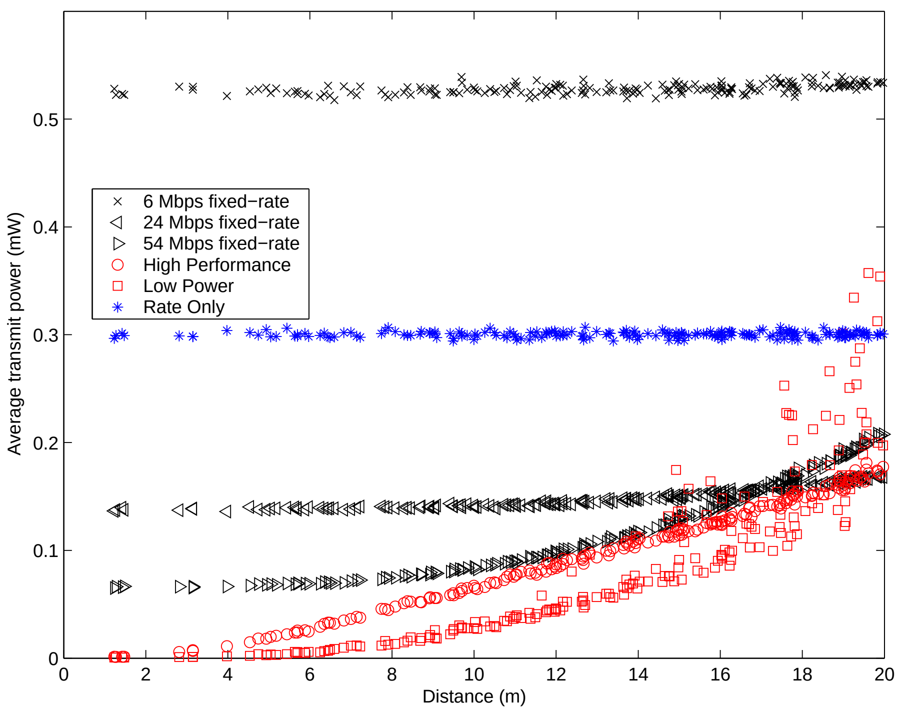

(b)

Figure 11: (a) Mean transfer times and (b) average transmit power in an 802.11a infrastructure WLAN with 10 WSTAs as function of the distance to the AP.

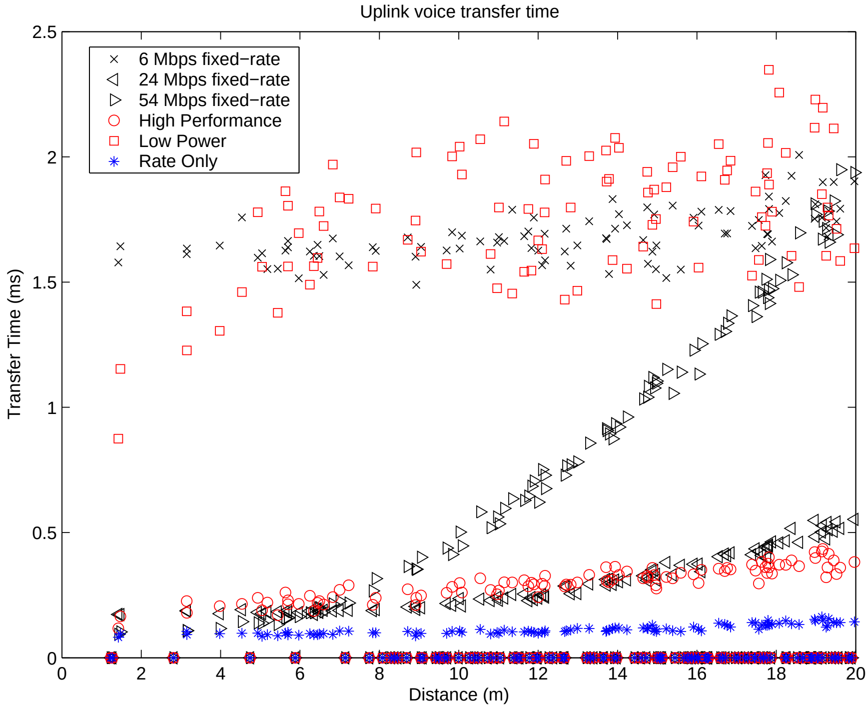

(a)

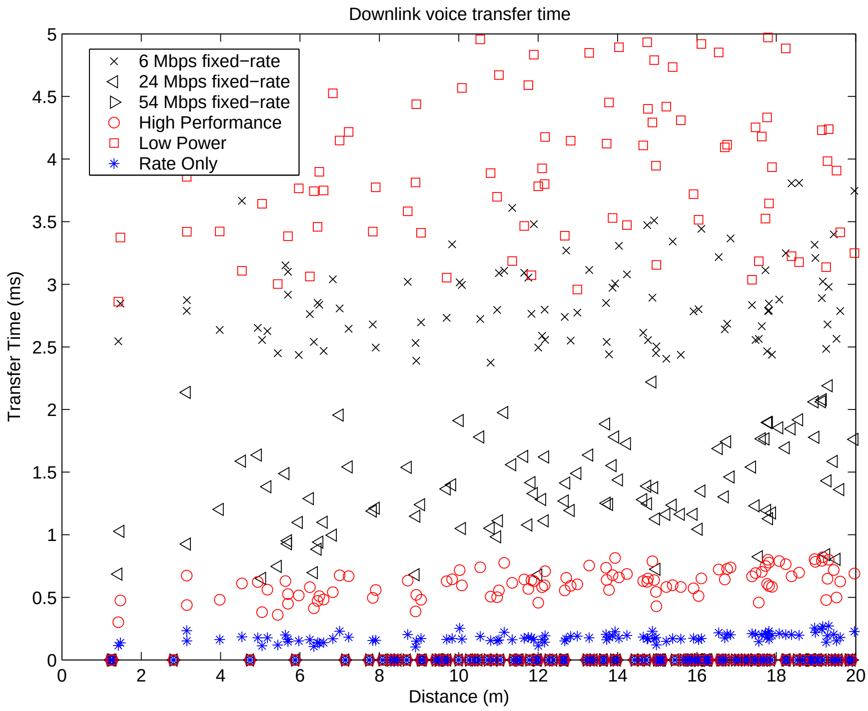

(b)

Figure 12: (a) Uplink and (b) downlink voice transfer times in an 802.11a infrastructure WLAN with 10 WSTAs as a function of the distance to the AP.

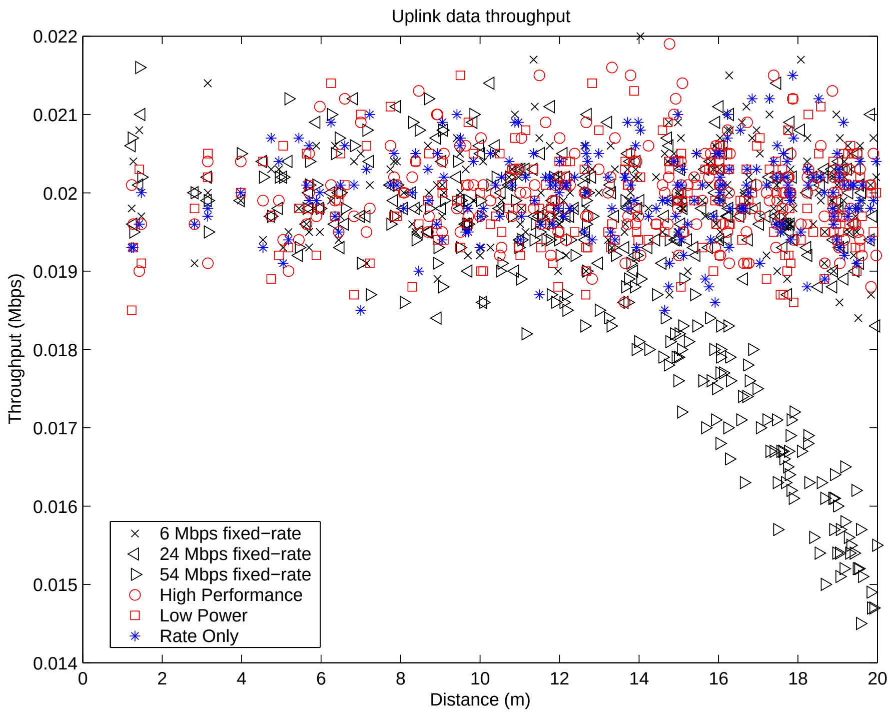

(a)

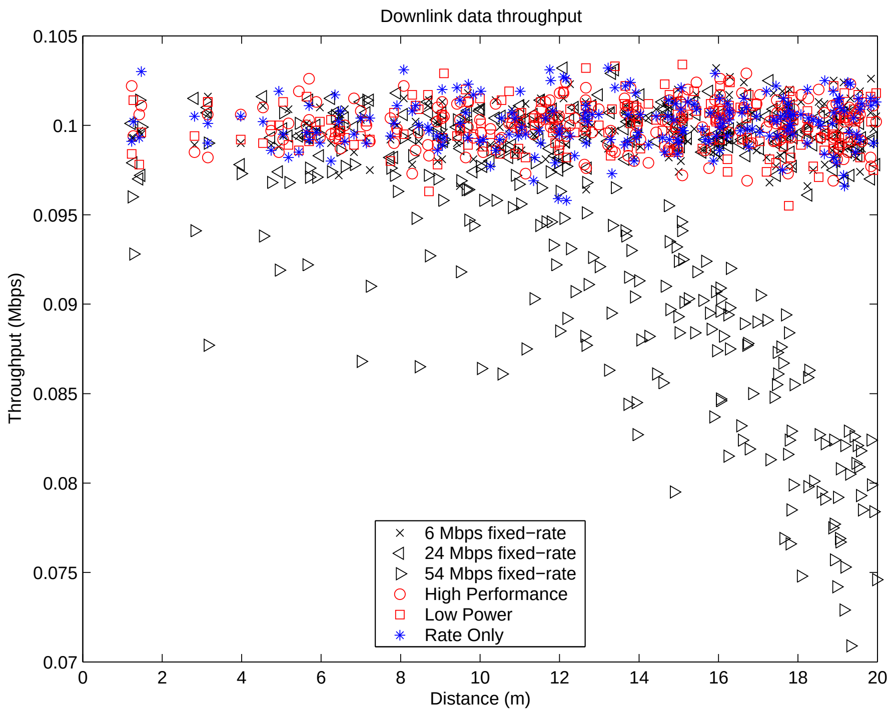

(b)

Figure 13: (a) Uplink and (b) downlink data throughput in an 802.11a infrastructure WLAN with 10 WSTAs as a function of the distance to the AP.

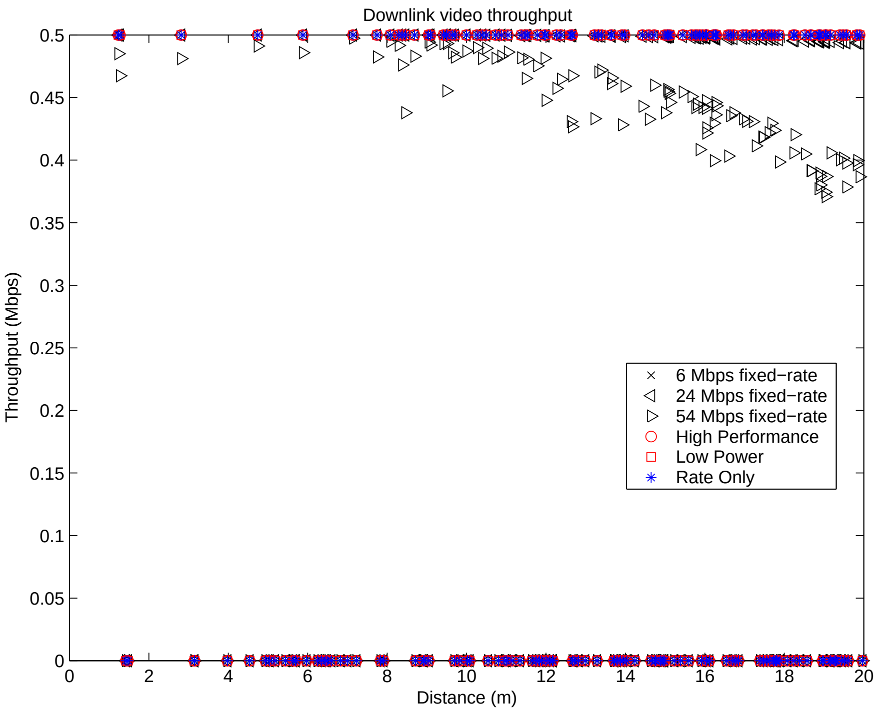

Figure 14: Downlink video throughput in an 802.11a infrastructure WLAN with 10 WSTAs as a function of the distance to the AP.

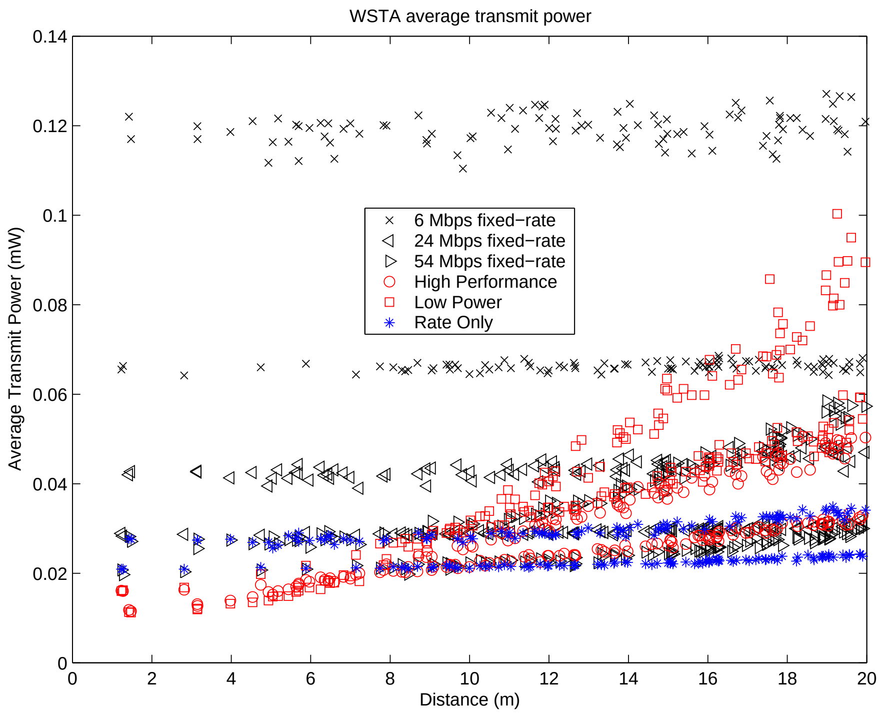

(a)

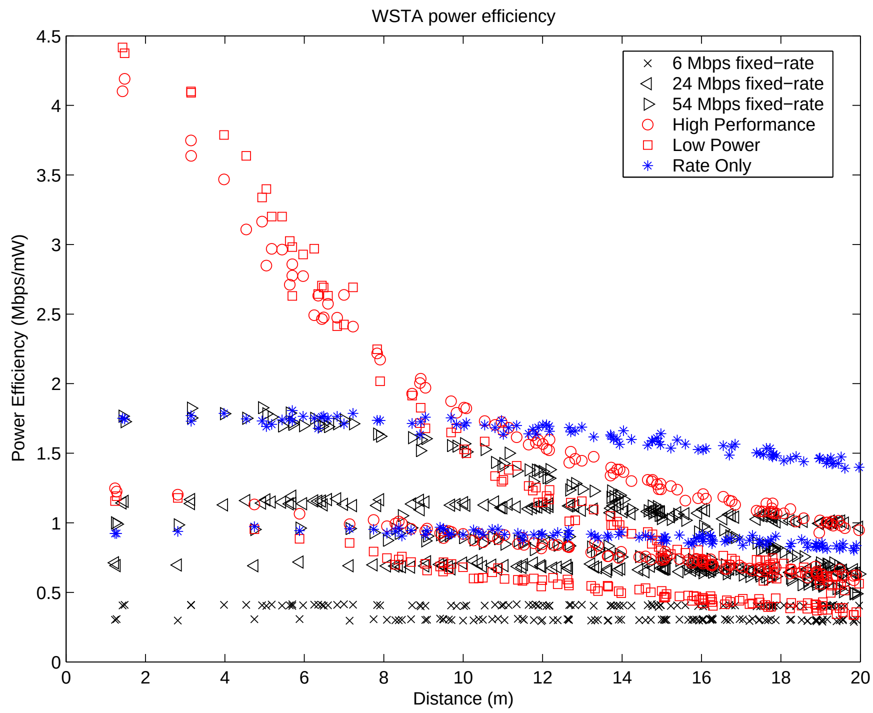

(b)

Figure 15: (a) Average transmit power and (b) power efficiency of a WSTA in an 802.11a infrastructure WLAN with 10 WSTAs as function of the distance to the AP.