This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the ICC 2008 proceedings.

GORA: Goodput Optimal Rate Adaptation for 802.11 using Medium Status Estimation

Nicola Baldo,† Federico Maguolo,† Simone Merlin,† Andrea Zanella,† Michele Zorzi,†,
Diego Melpignano‡, David Siorpaes ‡

† Department of Information Engineering, University of Padova, Italy
‡ ST-microelectronics, Agrate Brianza (MI), Italy
e-mail:{baldo,maguolof,merlo,zanella,zorzi}@dei.unipd.it
{diego.melpignano,david.siorpaes}@st.com

***Abstract*—Rate Adaptation for 802.11 has been deeply investigated in the past, but the problem of achieving optimal Rate Adaptation with respect not only to channel-related errors but also to contention-related issues (i.e., collisions and variations in medium access times) is still unsolved. In this paper we address this issue by proposing 1) a practical definition of the Medium Status in a multi–user 802.11 scenario in terms of channel errors, MAC collisions and packet service times, and a method for its estimation based on measurements; 2) an analytical model of the goodput performance as a function of the Medium Status; 3) a rate adaptation algorithm, called Goodput Optimal Rate Adaptation (GORA), which is based on this model. Unlike other Rate Adaptation schemes proposed in literature, which require either modifications to the IEEE 802.11 standard or cooperation among nodes, GORA is totally stand–alone and standard compliant. In fact, the Medium Status Estimation used by GORA is obtained by using standard MAC counters that are commonly collected by commercial MAC drivers, and no explicit interactions with the other devices in the network is required. Therefore, GORA offers the advantage of being readily deployable on real devices. The performance of GORA is evaluated through NS2 simulations which reveal that, as expected, GORA outperforms other well-known Rate Adaptation algorithms in several scenarios and can be used as a new reference benchmark.**

I. Introduction

In wireless systems, such as WLANs, the propagation environment changes over time and space due to factors such as mobility and interference, thus impacting on the channel reliability. To cope with this challenging environment, many wireless interfaces offer the possibility of dynamically tuning some system parameters in order to adapt to the environmental variations. The IEEE 802.11 specifications, in particular, define a plurality of PHY modes which can be used for the transmission of data frames. Each PHY mode uses a particular modulation and channel coding scheme and, consequently, offers different performance in terms of transmission duration, overhead, and robustness against reception noise and interference.

In recent years, there has been a significant amount of research on Rate Adaptation (RA) algorithms. In particular, the case of a single sender/receiver pair has been deeply investigated [1]–[6].

However, in typical 802.11 scenarios, multiple users compete for the medium. Therefore, in addition to channel-related packet losses, also MAC collisions and variable medium access time have a significant impact on performance, and in practice make the above mentioned schemes sub-optimal and, in some cases, very inefficient. Some recently proposed practical RA schemes [7], [8] address the problem of collision-related packet losses; however, to the best of our knowledge, no previous work formulated an optimal RA policy for multi-user 802.11 scenarios.

In this paper, we try to fill this gap by proposing i) a way to estimate the status of the transmission medium, i.e. channel propagation and medium congestion and ii) a novel RA algorithm, named Goodput Optimal Rate Adaptation (GORA) which uses such an estimation together with an analytical model of the goodput performance for 802.11.

The Medium Status is defined as the triplet $< \text{SNR}, P_{\text{coll}}, \xi >$, where SNR is the Signal to Noise Ratio at the receiver, i.e., at the Access Point, P_{coll} is the Collision Probability experienced by the mobile station (STA) and ξ is the average tick period, defined as the time between two successive decrements of the backoff counters. In order to make it possible to implement our scheme on real devices, we provide a method for estimating the Medium Status based on information which is commonly available on commercial devices in the form of 802.11 Management Information Base (MIB) counters,[1] with the exception of an event counter that we propose to add and that could be easily implemented by device manufacturers.

Hence, we develop a mathematical model to compute the expected throughput for a given Medium Status, as a function of the selected PHY rate. The Goodput Optimal Rate Adaptation (GORA) algorithm, then, simply selects the PHY mode that, according to the outcome of the analytical model, yields the better throughput for the estimated Medium Status. Unlike the current reference model for Rate Adaptation [1], GORA takes into account not only losses due to channel impairments, but also losses due to collisions and variations in the medium access time due to the backoff freeze procedure.

We use an enhanced version of the NS2 simulator to compare the performance of the most common RA schemes

[1]The formal specification of the 802.11 MIB is Annex D of the 802.11 specification. [9]

proposed in the literature with our proposal. Simulation results obtained by using perfect Medium Status estimation show that an ideal implementation of GORA with exact receiver SNR information always outperforms state–of–the–art RA algorithms in congested scenarios, thus establishing a new performance reference for 802.11 rate adaptation. Moreover, when using the SNR estimation provided by the proposed Medium Status Estimator, GORA achieves performance close to the ideal scheme, thus confirming that the RA framework here proposed is actually suitable for implementation.

The rest of the paper is organized as follows. In section II we analyze prior art RA algorithms, discussing their merits and drawbacks. In Section III we present the Medium Status Estimator and the analytical model for the throughput with different PHY rates, which is the core of GORA. Simulation results and comparisons with state-of-the-art algorithms are presented in Section IV. Finally, in section V, the conclusions are drawn.

II. Related Work

While the problem of 802.11 RA has been thoroughly investigated in the past, and many RA algorithms have been proposed in the literature, none of them has been successful in addressing all the issues which arise in a real 802.11 system. In particular, the most proposals aim at optimizing the PHY rate with respect to channel impairments only, neglecting the detrimental effect of collisions and variations in the medium access time.

In [1], the authors propose the MPDU-Based Link Adaptation Scheme (MBLAS), which makes use of an analytical model to evaluate the 802.11 goodput as a function of the SNR, the PHY mode and the payload size. The proposed model takes into account the 802.11 backoff and retransmission procedure, but is limited to the scenario with a single transmitter/receiver pair, for which MBLAS provides the theoretically optimal rate. It is to be noted that this is the current reference analytical model for 802.11 RA; however, this model is not valid in multi-user scenarios, and consequently, as we will show in this paper, MBLAS is subject to a significant degradation in performance as the number of users contending for the medium increases. Another issue is that the Received Signal Strength Indicator (RSSI) is used to determine the SNR and, consequently, to select the PHY mode. The RSSI, in fact, is a measure of the received signal power, which is proportional to the Signal to Noise Ratio (SNR) at the receiver. In practice, however, this approach is limited by a number of factors. For instance, it is a common experience that the RSSI returned by a wireless board circuitry is not always reliable. Furthermore, some schemes select the PHY mode according to the RSSI measured at the transmitter, assuming it is the same that would be experienced at the receiver. However, the assumption of symmetry is often not very accurate in reality.

In Receiver Based Auto Rate (RBAR) [2], the receiver STA selects the more suitable PHY mode on the basis of the RSSI measured during the reception of the RTS frame. The selected PHY mode is then communicated to the sender by using the CTS frame, so that the sender will adopot the chosen rate for the subsequent data transmission. While effective in overcoming channel asymmetry issues, this algorithm is not standard compliant since it requires modification to the RTS, CTS and data frame structure, as well as to the PLCP header, in order to include the necessary control information. Moreover, the proposed RSSI-based rate selection scheme at the receiver takes into account only the success probability of a single frame transmission, thus completely neglecting the impact of the MAC layer on the performance.

In [3] the authors propose a RSSI–based Link Adaptation strategy. The PHY mode is selected based on the measured RSSI, which is compared with dynamically defined thresholds. The use of dynamic thresholds aims at alleviating both the inaccuracy of RSSI measurements and the channel asymmetry issues. The drawback of this proposal is that the thresholds are adjusted considering the loss rate observed for a given PHY mode: this practice could easily lead, in case of frame losses due to collisions, to conservative threshold settings and, consequently, to lower throughput than what actually sustainable.

Due to the issues with RSSI measurements in real devices, discussed above, a completely different class of RA schemes has been developed which exploits the success history of past transmissions to infer the channel conditions. These schemes typically exhibit a much better implementability with respect to RSSI-based schemes, at the price of a certain degree of suboptimality.

The best known algorithm in this class is Auto Rate Fallback (ARF) [4] which is based on the following consideration. In the absence of interference from other users, a certain number of subsequent failures are likely due to a degradation of the SNR, so that a more robust rate has to be selected. Conversely, when a certain number of subsequent successful transmissions is observed, a higher rate is selected to improve throughput. These types of schemes are not subject to RSSI measurement inaccuracies nor to channel asymmetry issues. One of the drawbacks of ARF, however, is that it periodically tries a higher transmission rate[2] to check whether it is sustainable; this behaviour is inefficient in static scenarios where the optimal rate remains the same for prolonged periods.

Adaptive Auto Rate Fallback (AARF) [5] aims at alleviating this problem by applying a binary exponential backoff to the number of subsequent successful transmissions needed to try a higher rate. In this way, AARF is more stable than ARF and achieves better performance in static scenarios. Nonetheless, both ARF and AARF assume that packet losses are always due to channel errors, so that their performance can rapidly degrade in high traffic scenarios, where a significant amount of packet losses are caused by collisions.

Some other RA schemes try to combine the best features of the RSSI-based and loss-based approaches. For instance, Hybrid Rate Control (HRC) [6] exploits the measured RSSI and Frame Error Rate to distinguish between short-term and long-term variations of the channel conditions. This mechanism exploits a throughput-based rate controller which probes adjacent rates to determine if a rate switch is neccessary.

[2]As reported in [2], ARF tries a higher rate every 10 consecutive successful transmissions.

Moreover, two sets of thresholds (named stable and volatile low thresholds) are used depending on the detected variations of the RSSI. Again, this scheme does not consider the fact that packet losses might be also due to collisions.

To summarize, a major drawback of all the RA schemes discussed so far is that they are designed for scenarios in which a single node is transmitting on the wireless channel. In real situations, however, it is often the case that multiple nodes are competing for the medium. Consequently, due to the way the 802.11 MAC works (i.e., CSMA/CA with DCF), the goodput actually experienced by active nodes is influenced not only by channel-related packet losses, but also by MAC collisions and variations in the time required to access the medium. These factors cause the formerly discussed RA algorithms to achieve sub-optimal and, in some cases, very low performance. In particular, loss-based RA schemes such as ARF or AARF do not work properly, since losses due to MAC collisions can easily lead to the choice of a low-rate PHY mode even in cases in which a high rate is sustainable. As for RSSI-based schemes, the choice of RSSI thresholds is optimal only for the single user scenarios, but can easily become non-optimal as the time required for a successful packet transmission increases due to collisions and increased medium access delay.

More recently, some solutions have been proposed to address the problem of RA in congested 802.11 networks. For example, Closed Loop Adaptive Rate Allocation (CLARA) [8] is an ARF-like RA scheme which aims at reacting differently to losses due to channel errors and collisions, respectively. A significant drawback of this scheme is that it is based on the assumption that losses after a successful RTS/CTS exchange are always due to channel errors; consequently, it requires the use of the RTS/CTS handshake that has a significant cost in terms of overhead. The Collision Aware Rate Adaptation scheme proposed in [7] exploits the same mechanism for loss differentiation but implements an adaptive RTS/CTS probing scheme which reduces the overall RTS/CTS usage, thus somehow mitigating the inefficiency of CLARA. We note, however, that both CLARA and CARA do not consider the impact on the performance of the variations in the medium access time.

To conclude, no previous work, to the best of our knowledge, has provided a RA scheme which is optimal with respect to both channel impairments and contention-related issues, comprehensive of both frame collision probability and medium access times. The main reason for this is that no analytical models for the performance of 802.11 rate adaptation in multi-user scenarios have been presented so far. In the next section we propose such a model, which enables the definition of our Goodput-Otimal Rate Adaptation (GORA).

III. System Model

A. Medium Status Estimation

In order to develop a generic model for 802.11 performance, we are interested in differentiating between packet losses due to interference and packet losses due to channel impairment. For a given transmission, we denote the collision and channel error probability as P_{coll} and P_{err}, respectively.

The Medium Status is defined as the triplet $<\ \text{SNR}, P_{\text{coll}}, \xi\ >$, where SNR is the Signal to Noise Ratio at the receiver, P_{coll} is the Collision Probability experienced by the mobile station (STA) and ξ is the average *tick period*, defined as the time between two successive decrements of the backoff counters, thus accounting for the freezing of the backoff counter when the medium is sensed busy [10].[3]

This Medium Status definition permits to take into account all PHY and MAC layer aspects which contribute to the maximum goodput achievable by a 802.11 STA. For this reason, in order for an RA scheme to be effective on real devices, it is of primary importance for a STA to be able to determine the current Medium Status.

We propose a Medium Status Estimation method which is based on the use of some measurements available at the MAC layer of the STA performing RA. These measurements, and the notation which will be used for them throughout this paper, are:

- t_s: the number of successfully transmitted unicast MSDUs, i.e., the number of transmitted data frames for which an ACK was received;
- t_f: the number of transmitted data frames for which an ACK was not received;
- r_s: the number of successfully received data frames, including those not addressed to the STA being considered;
- r_f: the number of received frames for which the Frame Check Sum (FCS) failed;
- s_i: the number of idle time–slots, i.e., 802.11 PHY slots in which the channel was sensed idle, excluding those belonging to an inter- or intra-frame space.

All these measurements can be obtained directly or indirectly by some of the counters available within the 802.11 Management Information Block (MIB).[4] The only exception is the idle time–slot counter s_i, which is not listed among the counters in the MIB; we note, however, that its implementation would be rather straightforward, and therefore our proposal still maintains a high degree of implementability in real devices. We assume that all above mentioned counters refer to the events occurred in a time window of given duration D.

We consider P_{coll} as the probablity that a packet is erroneously received due to interference at the receiver. With this definition, P_{coll} is hard to estimate in a general interference environment; consequently, we approximate it with the probability that a given transmission is simultaneous with at least one interfering transmission within a given interference range of the transmitter. This approximation is accurate when all terminals are within carrier sense range, whereas it becomes less accurate in the presence of hidden/exposed terminals.

Following the same approach proposed in [11], we assume that both the transmission and the collision probability are stationary, i.e., independent of the particular slot considered. With

[3] In [11] this is called *slot period*. We prefer to use the term *tick period* in order to avoid possible confusion with the 802.11 PHY slot time – the difference is that the tick period has random duration due to the backoff freeze procedure, and its average value can be much longer than the 802.11 PHY slot time.

[4] We refer to the `dot11Counters` described in the IEEE 802.11 standard, Annex D [9]. We note that, in order to derive the measurements we need from the MIB counters, some processing is required, since some counters also include control and management frames, while the measurements we use are supposed to only check for data frames.

this assumption, the collision probability for a transmission by the STA under consideration equals the probability that a randomly selected time slot is occupied by a transmission by another STA. This yields:

$$P_{\text{coll}} = \frac{r_s + r_f}{r_s + r_f + s_i}. \quad (1)$$

The resulting P_{coll} does not depend on the transmissions by the STA under consideration, but only on the events caused by all the other STAs. This is the main difference with the estimator presented in [12], which allows our estimator to be effective also in the presence of packets losses due to channel impairments.

The packet loss probability P_{loss} can be obtained as:

$$P_{\text{loss}} = \frac{t_f}{t_f + t_s}. \quad (2)$$

We note that P_{loss} accounts for the contribution of both P_{coll} and P_{err}. Since errors due to noise and interference are overlapping and independent events, we have

$$P_{\text{loss}} = P_{\text{coll}} + P_{\text{err}} - P_{\text{coll}} P_{\text{err}} \;, \quad (3)$$

and combining (2) and (3) we get

$$P_{\text{err}} = \frac{t_f - (t_f + t_s) P_{\text{coll}}}{(t_f + t_s)(1 - P_{\text{coll}})} \;. \quad (4)$$

We suppose that P_{err} is univocally determined by a known function of the PHY mode being used, the packet size and the SNR (without considering the effect of interference) seen by the receiver, and that the SNR is constant for the whole packet transmission duration. Then we determine SNR by inverting the SNR versus P_{err} relationship for the rate being used. Clearly this practice requires that the same PHY mode was used for the whole observation period D. This method is useful to get an SNR estimation without using the RSSI measurement, which, as discussed in Section II, is not reliable in real devices.

Finally, the average tick period ξ can be estimated dividing the observation period D by the total number of tick events (i.e., idle slots and busy periods) counted in the observation window.

B. Goodput Model

We define the service time y of a MAC frame (MPDU) as the time since the MPDU was scheduled for transmission over the air interface until either the corresponding ACK frame is correctly received by the sender or the MPDU itself is discarded due to exceeded retransmission limit.

Clearly, for each MPDU, y is a random variable. We assume that the STA always has a frame to transmit. In this case, the service time of a frame can be modelled as a stochastic process that renews itself after each service. Let L be the payload size of a MPDU, which we assume to be constant. The goodput G of the system can be expressed as

$$E[G] = \frac{L}{E[y]} \left(1 - (P_{\text{loss}})^{r_{\text{max}}}\right) \;; \quad (5)$$

where r_{max} is the max retry limit, i.e., the maximum number of transmission attempts a MPDU can undergo before being dropped by the MAC layer.

Let b denote the cumulative number of backoff decrement events occurred during the service of the MPDU being considered. Let T_{PDU} denote the total time spent transmitting and retransmitting the MPDU being considered.

We can express the average service time as

$$E[y] = \xi E[b] + E[T_{\text{PDU}}] \;. \quad (6)$$

The average time in the backoff is given by

$$E[b] = \sum_{i=1}^{r_{\text{max}}} (P_{\text{loss}})^{i-1} (1 - P_{\text{loss}}) \sum_{j=1}^{i} \left(\frac{CW(j) - 1}{2}\right) + (P_{\text{loss}})^{r_{\text{max}}} \sum_{k=1}^{r_{\text{max}}} \frac{CW(k) - 1}{2},$$

where $CW(j)$ is the size of the contention window at the j–th transmission attempt of the same MPDU (see [10], [11]). The average transmission time is given by

$$\begin{aligned} E[T_{PDU}] &= \sum_{i=1}^{r_{\text{max}}} (P_{\text{loss}})^{i-1} (1 - P_{\text{loss}}) ((i-1)T_f + T_s) \\ &+ (P_{\text{loss}})^{r_{\text{max}}} r_{\text{max}} T_f = \\ &= \frac{1 - r_{max} P_{loss}^{r_{max}-1} + (r_{max} - 1) P_{loss}^{r_{max}}}{1 - P_{loss}} P_{loss} T_f \\ &+ (1 - P_{loss}^{r_{max}}) T_s + (P_{\text{loss}})^{r_{\text{max}}} r_{\text{max}} T_f \;; \end{aligned}$$

where T_s denotes the average transmission time of a MPDU, which includes the transmission time, the short inter–frame space (SIFS), the ACK transmission time and the distributed inter–frame space (DIFS), whereas T_f is the time spent in case of a failed transmission, including the extended inter–frame space (EIFS).

C. Rate Adaptation

The proposed rate adapatation algorithm uses the outcome of Medium Status estimation $< \text{SNR}, P_{\text{coll}}, \xi >$ to compute the expected goodput for all the possible PHY modes by means of (5). The PHY mode achieving the highest goodput is then selected.

The optimization is performed periodically every T_{opt} seconds. As the Medium Status is estimated at run time by collecting MAC counters statistics the duration of T_{opt} is related to D. We note that both values should be chosen taking into account the desired tradeoff between an accurate estimation and a fast PHY mode adaptation.

IV. Performance Evaluation

A. Simulator

Simulations have been performed by using an enhanced version of the ns2 simulator [13]. All the specific parameters for the network setting are compliant with the IEEE 802.11g standard. In particular, unless otherwise specified, r_{max} has been set to 7. The interference model used in the simulations is based on a Gaussian approximation for the interference, as widely assumed in literature. The propagation model is determined by the two ray ground model accounting only for the path loss component. Fast channel gain fluctuations are not considered.

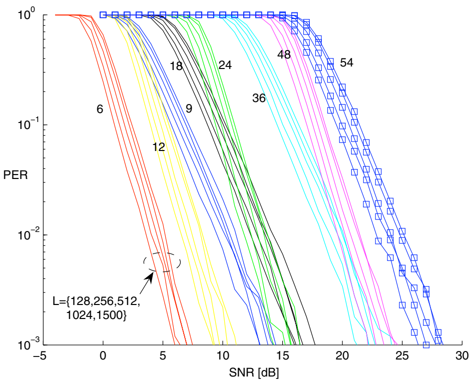

Figure 1. P_{err} vs SNR for different 802.11g PHY modes and MPDU sizes

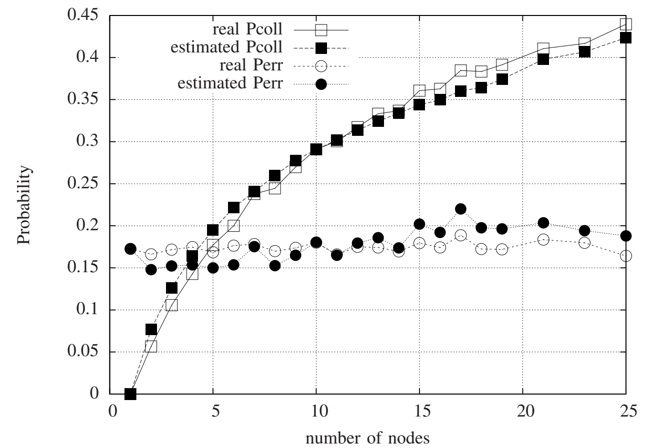

Figure 2. Actual and estimated P_{coll} and P_{err} as a function of the total number of nodes; the test node has SNR $= 20.54$ dB

We use a stochastic error model for PHY layer transmissions, according to which packet losses are independent and occur with a probability which is determined as a function of the PHY mode, Signal to Interference plus Noise Ratio (SINR), and packet length. This packet error probability has been computed offline using a dedicated PHY layer simulator accounting for the standard specifications of the OFDM modulation and coding in 802.11g. The curves are reported in Fig. 1.

The considered scenario refers to an infrastructure network where a test STA is connected to the access point (AP) and is provided with the parameter estimation and rate adaptation algorithms. Other STAs are connected to the same access point in order to simulate an interfered scenario. To better evaluate GORA performance the interfering STAs mantain a fixed transmission rate are not provided with the rate adaptation algorithms.

Only uplink connections are simulated. All the traffic is UDP with a packet length of 1500 bytes; the packet generation rate is high enough so that all STAs always have pending transmissions.

B. Estimation validation

In this section, the proposed estimation technique for P_{coll} and P_{err} is validated.

In Fig. 2 we report the results for a scenario where the test node is placed at a fixed distance from the AP, and an increasing number of interfering nodes are placed close to the AP. All nodes use the 54 Mbps mode. The value of P_{coll} and P_{err} estimatioed according to formulae (1) and (4) are compared with the actual values measured from the trace files. As it can be seen, estimated values closely match the actual ones. We obtained similar accuracy in a wide range of SNR and number of interfering nodes, and we can therefore conclude that the proposed estimator is able to provide us accurate information on the propagation environment and congestion level. In the case of very high SNR, i.e., no channel errors are present, the estimated P_{coll} is very close to the one provided by the analytical model in [11].[5] Due to space constraints, additional validation results are not reported.

[5]For these simulations we used $r_{\mathrm{max}} = 100$ in order to match with the infinite retransmission limit assumption in [11].

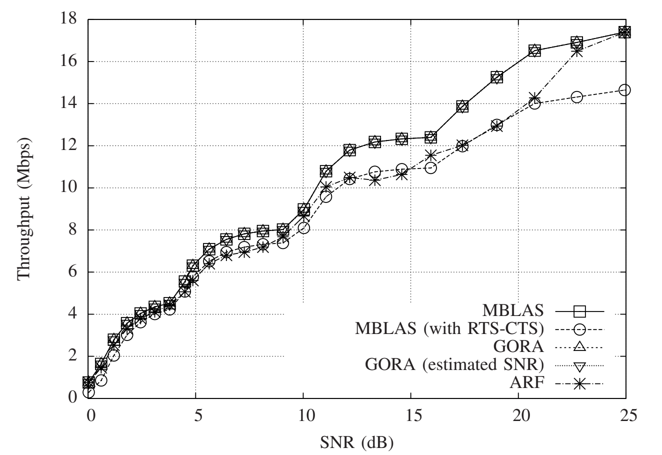

Figure 3. Throughput of rate adaptation algorithms, as a function of the SNR, in case of no interference.

C. Rate adaptation performance

In this section we report the performance evaluation of the proposed rate adaptation mechanism. As we discussed in Section II, many existing RA schemes are based on the RSSI measurement, which is not reliable on real devices; consequently, these schemes have very good performance in ideal conditions, but poor implementability in practice. To provide a fair comparison, we test the proposed GORA algorithm both in the case an exact SNR value is used and in the case the SNR value is provided by the SNR estimator discussed in Section III-A. These two versions of GORA are compared with ARF, MBLAS (which assumes to have perfect SNR knowledge at the receiver) and a modified version of MBLAS which uses the RTS-CTS mechanism to receive information on the SNR at the receiver.[6]

In Figure 3 we report the results for a scenario with no interferers. In this case, MBLAS has been shown to achieve optimum performance. The proposed GORA algorithm, both with

[6]This combines the MBLAS model with the information exchange protocol of RBAR. The former provides enhanced performance, while the latter offers enhanced implementability since it can cope effectively with channel asymmetries.

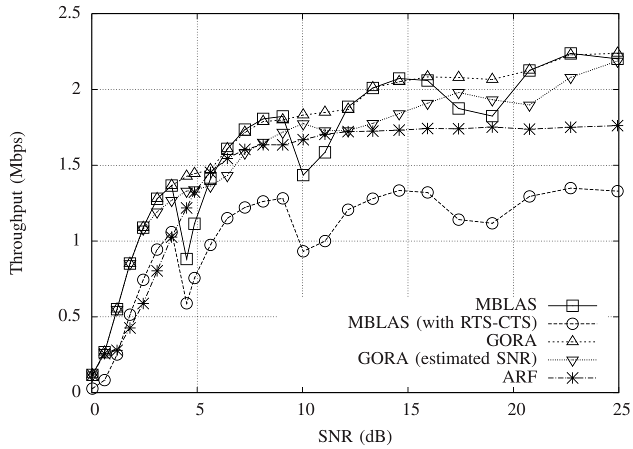

Figure 4. Throughput of rate adaptation algorithms, as a function of the SNR, in presence of 9 interfering nodes.

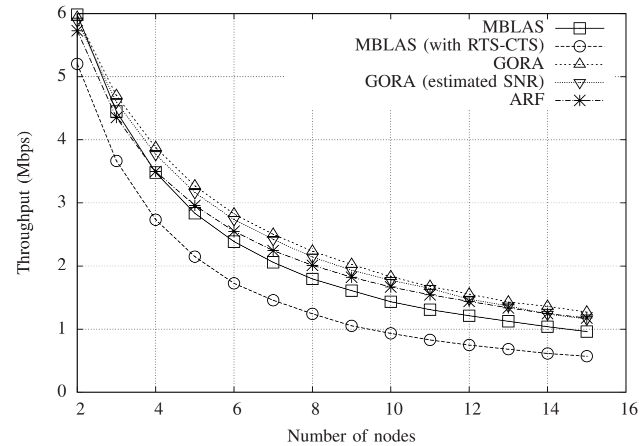

Figure 5. Throughput of rate adaptation algorithms, as a function of the total number of nodes; the test node has SNR $= 10\,$dB.

exact and estimated SNR, achieves the same optimum throughput. These algorithms outperform ARF and the MBLAS version, which suffer from the RTS-CTS overhead.

In Figure 4 we report the results obtained for a scenario with 9 interferers. In this case, MBLAS turns out to be suboptimal. This is due to the fact that MBLAS does not account for the losses due to collisions, nor for the increased medium access time due to other users accessing the channel. Instead, the proposed GORA mechanism is able to adapt the PHY mode selection to the medium status and, in case an exact SNR knowledge is considered, it always achieves the highest throughput. We note that, unlike MBLAS, GORA with exact SNR knowledge takes into account all factors that determine the performance, and hence performs an optimal RA, and is therefore a candidate as the new performance benchmark for 802.11 RA. When using the estimated SNR value, GORA still shows good performance. In particular we stress that this version of GORA is very suitable for implementation on real network devices. In this regard we are interested in the comparison with ARF, which is widely deployed in practice, and with MBLAS with RTS-CTS, which is a rather practical (though not standard-compliant) implementation. Both algorithms are outperformed by GORA.

In Figure 5 we report the results for a fixed SNR and a variable number of interferers. In this particular setting, GORA both with exact and estimated SNR outperforms the other rate adaptation algorithms.

V. Conclusions

In this paper we have proposed a Goodput Optimal Rate Adaptation (GORA) scheme for IEEE 802.11g. This algorithm is based on the estimation of the Medium Status, which accounts for collision probability, SNR at the receiver and medium occupancy time. The estimation process uses the 802.11 MAC MIB counters provided by the standard, plus an additional counter which we identify to be of fundamental importance and which could be easily implemented by device manufacturers. On top of this Medium Status Estimation, we developed an analytical model for the 802.11 throughput performance, which we propose as the new reference model for 802.11 RA.

Simulations showed that the GORA algorithm using perfect SNR information outperforms RA schemes previously regarded as optimal; at the same time, when using the proposed SNR estimator, GORA was shown to outperform other state-of-the-art algorithms having similar or lower degree of implementability.

References

[1] D. Qiao, S. Choi, and K.G. Shin. Goodput analysis and link adaptation for IEEE 802.11 a wireless LANs. *IEEE Transactions on Mobile Computing*, 1(4):278–292, 2002.

[2] G. Holland, N. Vaidya, and P. Bahl. A rate-adaptive MAC protocol for multi-Hop wireless networks. *Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 236–251, 2001.

[3] J. del Prado Pavon and S. Choi. Link adaptation strategy for IEEE 802.11 WLAN via received signal strength measurement. *IEEE International Conference on Communications*, pages 1108–1113, 2003.

[4] A. Kamerman and L. Monteban. WaveLAN-II: A high-performance wireless LAN for the unlicensed band: Wireless. *Bell Labs technical journal*, 2(3):118–133, 1997.

[5] M. Lacage, M.H. Manshaei, and T. Turletti. IEEE 802.11 rate adaptation: a practical approach. *Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, pages 126–134, 2004.

[6] I. Haratcherev, K. Langendoen, R. Lagendijk, and H. Sips. Hybrid Rate Control for IEEE 802.11. *ACM International Workshop on Mobility Management and Wireless Access Protocols (MobiWac)*, pages 10–18.

[7] J. Kim, S. Kim, S. Choi, and D. Qiao. CARA: Collision-Aware Rate Adaptation for IEEE 802.11 WLANs. *IEEE INFOCOM*, 2006.

[8] C. Hoffmann, MH Manshaei, and T. Turletti. CLARA: Closed-Loop Adaptive Rate Allocation for IEEE 802.11 Wireless LANs. In *International Conference on Wireless Networks, Communications and Mobile Computing*, volume 1, 2005.

[9] *ANSI/IEEE Std 802.11, 1999 edition*.

[10] Andrea Zanella and Francesco De Pellegrini. Statistical characterization of the service time in saturated IEEE 802.11 networks. *IEEE Communications Letters*, 9(3):225–227, 2005.

[11] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, 2000.

[12] G. Bianchi and I. Tinnirello. Kalman filter estimation of the number of competing terminals in an IEEE 802.11 network. *Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE INFOCOM 2003)*, 2, 2003.

[13] N. Baldo, F. Maguolo, M. Miozzo, M. Rossi, and M. Zorzi. Ns2-miracle: a modular framework for multi-technology and cross-layer support in network simulator 2. In *ACM NSTools*, 2007.