

Impact of Complex Wireless Environments on Rate Adaptation Algorithms

Defeng Xu
 Mobile Systems Lab
 Dept. of Computer Science
 University of California, Los Angeles
 Email: dxu@cs.ucla.edu

Rajive Bagrodia
 Mobile Systems Lab
 Dept. of Computer Science
 University of California, Los Angeles
 Email: rajive@cs.ucla.edu

Abstract—Due to the popularity of mobile computing devices, 802.11-based wireless networks have become common place. A critical but unspecified component of 802.11 is transmission rate adaptation. Rate adaptation is difficult to do because it requires accurately estimating channel conditions while experiencing distinctly different wireless phenomena such as multi-path fast fading, nodal mobility, environmental mobility and hidden terminal interference. Rate adaptation algorithms use various methods to estimate channel conditions but these methods can be inaccurate in complex wireless environments causing incorrect adaptations. In this paper, we investigate the impact of inaccurate channel estimation methods by studying two rate adaptation algorithms, CHARM and SampleRate. We emulate high fidelity versions of both algorithms, using QualNet and evaluate their performance under environmental mobility and hidden terminal interference conditions. We found that the performance of both algorithms degrade dramatically due to inaccurate estimations caused by stale channel information. We propose two optimization techniques, REGS SINR and IntRec, to counteract these inaccuracies and implement them for CHARM and SampleRate. The proposed optimization techniques increase the performance dramatically for CHARM, up to 340% for environmental mobility and more than 290% for hidden terminal interference. We saw similar improvements of 220% and 88% for SampleRate.

Index Terms—rate adaptation, 802.11, CHARM, SampleRate, adaptation, 802.11, CHARM, SampleRate

I. INTRODUCTION

Mobile computing devices have become ubiquitous. From standard laptops to cutting edge smart phones, people are no longer tethered by their desktop PCs for their computing needs. One of the most common ways for these mobile devices to connect to each other and to the Internet is through IEEE 802.11-based wireless networks. A critical component of 802.11-based wireless networks is transmission rate adaptation, which is currently left unspecified in the standard itself [7]. The fundamental trade off in transmission rate selection is choosing between higher data rate and signal robustness [4]. While a higher data rate can increase throughput and reduce the transmission time of packets, packets also have a lower chance to be successfully decoded at lower SINR. Therefore a rate adaptation algorithm needs to select the highest rate so that packets can be reliably decoded for the given channel conditions. What makes rate adaptation difficult to do is accurately estimating channel conditions while experiencing

distinctly different wireless phenomena such as multi-path fast fading, nodal mobility, environmental mobility and hidden terminal interference.

Current rate adaptation algorithms use various methods to estimate channel conditions such as SINR measurements, consecutive successes/losses, and transmission times [7]. However most of these methods assume the wireless environment is composed of mainly multi-path fast fading and nodal mobility effects. In more complex and dynamic environments, with additional effects such as environmental mobility and hidden terminal interference, these methods can give an inaccurate view of the wireless channel. This inaccuracy will cause the rate adaptation algorithms to adapt incorrectly and consequently degrade the network's performance.

In this paper we investigate the impact of environmental mobility and hidden terminal interference effects on channel estimation methods by studying two popular rate adaptation algorithms: CHARM [4] and SampleRate [1]. For this purpose, we emulate high fidelity versions of the two algorithms using QualNet [2], an industry standard network simulator, that allows us to keep the evaluation environment as realistic as possible while leveraging QualNet's rich suite of wireless simulation tools.

Our experiments found that the performance of both CHARM and SampleRate degrade dramatically due to the use of stale channel information, up to 77% for CHARM and 66% for SampleRate. Due to the 2-state nature of environmental mobility [5][6] and hidden terminal interference, information gathered when the channel is in one state may not be applied to the other. In the case of CHARM, channel measurements cannot be collected without a correctly decoded packet. CHARM's use of SMA to aggregate previous SINR measurements means for a lost packet it will use old measurements to determine current channel conditions. These measurements can be stale if they were taken during another state that does not relate to the current one. For SampleRate, a transmission rate eliminated from consideration due to four consecutive lost ACKS in one state does not mean it will not succeed in the other.

To counteract these performance degradations we introduce two optimization techniques that try to recover from the effects of the stale information caused by environmental mobility

and hidden terminal interference and implement them for CHARM and SampleRate. These two techniques increase the performance dramatically for CHARM, up to 340% for environmental mobility and more than 290% for hidden terminal interference. We saw similar improvements of 220% and 88% for SampleRate.

The rest of the paper is organized as follows. Section 2 provides brief overviews of CHARM, SampleRate and environmental mobility. Section 3 discusses how inaccurate channel estimations can occur. Section 4 shows the impact of inaccurate channel estimations on CHARM and SampleRate. In section 5 we propose our optimization techniques. Section 6 shows the performance impact of the proposed optimizations and section 7 concludes this paper.

II. BACKGROUND

In this section we will give a brief overview of the two rate adaptation algorithms we consider, CHARM and SampleRate, and the wireless phenomena environmental mobility.

A. CHARM

CHARM is a state of the art rate adaptation algorithm targeted for mobile users of 802.11. First proposed in a paper by Judd et al [4], CHARM is composed of three main components, receiver SINR estimation, rate selection, and on-line threshold adaptation.

SINR estimation - Receiver SINR estimation is used by CHARM to estimate the quality of the wireless channel at the intended destination of a new transmission. Noise and transmission strength at the receiver will be provided explicitly and piggybacked onto normal packets. The transmitter will measure the received signal strength of packets coming from the receiver and calculate the path loss. Using its own transmission power and the receiver noise the receiver SINR can be estimated. CHARM will do this continuously using a time-based SMA to aggregate the estimations.

Rate selection - CHARM uses a multi-rate retry mechanism to do rate selection. For each new packet CHARM will select the highest rate that has a threshold below the estimated receiver SINR. After a specified amount of consecutive failures CHARM will quickly jump to the next lowest rate.

On-line threshold adaptation - The reason CHARM adapts the rate thresholds over time is because the success chance of a rate at a certain SINR value will change depending on imperfections in transmit power information, receiver noise estimation, unreported interference, and multipath effects [4]. Given initial SINR thresholds for each rate, CHARM will also adapt the thresholds through an on-line process that occurs every few seconds. CHARM records the successes and failures of each transmission and at what predicted SINR those transmissions occurred at. These records are put into buckets based on the difference between the estimated SINR and the threshold SINR and whether it failed or succeeded. CHARM collects these records and uses them to make a decision on how to adapt the thresholds.

B. SampleRate

SampleRate is a widely studied rate adaptation algorithm originally proposed by Bicket [1]. The main feature of this algorithm is the use of average transmission times to determine rate selection. This algorithm is composed of two main components, transmission time collection and rate selection.

Transmission time collection - For each packet successfully transmitted SampleRate records the rate and the required number of retransmissions, which it then uses to calculate the transmission time. To avoid use of stale information, the average transmission time is then calculated over an averaging window of 10 seconds.

Rate selection - The rate is selected following three basic rules. The first rule is that any rate with four consecutive missed ACKs will be disqualified for that averaging window. To sample higher rates, if the packet is the 10th packet transmitted then SampleRate will select a random rate with a lossless transmission time lower than the current rate's average transmission time. Otherwise, SampleRate selects the rate with the lowest average transmission time.

C. Environmental Mobility

Environmental mobility (EM) refers to the ambient motion of objects within the vicinity of wireless devices [5][6]. When a mobile object obstructs the LOS between the transmitter and the receiver it will block signal ray clusters of the transmitter which causes significant signal loss creating a distinctly different channel condition. Therefore the wireless channel under EM exists under two distinct states, an unobstructed "good state" and a shadowed "bad state". State transitions are induced by mobile objects crossing the LOS between the transmitter and receiver. The wireless channel under the shadowed state can be modeled as a combination of diffraction loss due to the obstructing object and multi-path fast fading loss of the resulting diffracted signal.

If the mobile objects are people, the duration of the shadowed state is an order of magnitude greater than the duration of multi-path fast fading (i.e. 200ms compared to 20ms). The magnitude of the effect is also much greater, as signal strength degrades up to 20dB with just two people crossing the LOS [6]. These differences demonstrate that environment mobility effects are uniquely different from multi-path fast fading effects and protocols that are tailored to counteract multi-path fast fading and path loss effects will react differently under EM.

D. Hidden Terminal Interference

Hidden terminal interference occurs when there is a third terminal that is within range of a receiver but not the transmitter [3]. Since the third terminal is "hidden" from the transmitter, CSMA will not prevent collisions from occurring at the receiver. Hidden terminal interference is difficult for rate adaptation algorithms to adapt to because the standard procedure of lowering the transmission rate in response to packet loss will actually hinder overall performance if the hidden terminal is unresponsive. Similarly to EM, interference

also exhibits a 2-state nature, where the "good state" is when the channel is clear and the "bad state" is when it is not.

III. INACCURATE CHANNEL ESTIMATIONS

The purpose of a rate adaptation algorithm is to select the highest rate to transmit a packet that can be reliably decoded for the current channel conditions. To achieve this, various methods are used to estimate channel conditions such as SINR measurements, consecutive successes/losses, and transmission times [7]. If these methods assume a wireless environment composed mainly of multi-path fast fading and nodal mobility effects they will give inaccurate estimations in more dynamic and complex environments that also have EM and hidden terminal interference effects. This inaccuracy is mainly due to the unique binary state nature of EM and interference effects.

We look in particular at how these effects impact the estimation methods used by CHARM and SampleRate. The choice of these two algorithms was partly dictated by their use of differing techniques for channel estimation, where CHARM uses SINR measurements combined with rate fallback, and SampleRate uses a combination of transmission times and consecutive losses.

A. CHARM

CHARM's method of channel estimation uses SINR measurements which require a packet to be successfully decoded at the link layer. Since rate adaptation is done per connection basis, even a successful signal lock at the physical layer does not suffice. This is not significant if the main cause of loss is due to path loss from nodal mobility. Since path loss is a gradual and continuous process, previous measurements should still accurately reflect current channel conditions. If the main cause of loss is due to multi-path fast fading, because the SINR is already relatively low for fading effects to generate significant losses, the previous channel measurements should properly reflect this. However in the case of EM and interference, the channel does not degrade in a continuous manner, and when it does degrade, it degrades severely. Therefore the previous SINR measurements likely have no relevance to current channel conditions.

B. SampleRate

SampleRate's method of channel estimation uses consecutive losses to disqualify a rate from consideration. In the case of EM the channel exists in a binary state where the information collected during the "bad state" has no relevance to the "good state". Therefore disqualifying a rate due to losses that occurred during a "bad state" means that rate cannot be used even when the channel transitions back to a "good state." Since the duration of the averaging window is an order of magnitude greater than EM state durations, an incorrect disqualification will greatly hinder performance. Similarly disqualifying a rate due to losses caused by interference means that rate cannot be used when the channel is in the clear, "good state". The impact is not seen for nodal mobility or multi-path fast fading since there no "good state" for the channel to transition back to.




Fig. 1. Average throughput of 802.11b transmission rates at various transmitter/receiver distances

IV. PERFORMANCE IMPACT

In this section we investigate the impact of inaccurate channel estimations for CHARM and SampleRate. We first discuss our simulation environment then we evaluate the performance of the two algorithms under two wireless scenarios, EM and hidden terminal interference. The results show that the performances of both CHARM and SampleRate degrade significantly under EM and hidden terminal interference.

A. Simulation Environment

We implement a high fidelity emulation of CHARM and SampleRate in QualNet based on the specifications from the original paper, and in the case of CHARM we use the source code of the original implementation. Our simulation environment uses 802.11b which has four transmission rates, 11Mbps, 5.5Mbps, 2Mbps, and 1Mbps. For our radio model we use Ricean fading with two ray path loss. We use a CBR application that uses the UDP transport protocol to send items of 40kB every 0.04s which generates a client throughput of 8Mbps.

To determine how the different transmission rates perform in our simulation environment we ran multiple experiments with the transmitter and receiver at various distances apart using different fixed rates. The results of these experiments are shown in figure 1. From these results we determine heuristically what the optimal transmission rates are for a given distance. We evaluate the impact of inaccurate channel estimations for CHARM by comparing the throughput with on-line threshold adaptation turned on and off. In the case of SampleRate we turn on and off the use of four consecutive losses to disqualify a rate.

B. Environmental Mobility

In this scenario we consider the effects of EM [6][5], a common occurrence in real world wireless environments such as a WiFi AP located inside a crowded campus hallway. Our scenario is setup to model this crowded hallway environment. The distance between the transmitter and receiver are ranged from 100m to 200m which varies the optimal rate from 11Mbps to 1Mbps. For EM we use three movement profiles with the probability of transitioning to a shadowed state set to 30%, 50%, and 70% with an average state duration of 100ms.




Fig. 2. CHARM average throughput at various distances with on-line threshold adaptation on/off under EM using 30, 50, 70 % shadowed state




Fig. 3. SampleRate average throughput at various distances with consecutive losses on/off under EM using 30, 50, 70 % shadowed state

The signal loss due to the shadowed state is set to an average of 25dB which corresponds to 2-3 people crossing the LOS [6]. The results show significant performance degradations which can be seen for CHARM in figure 2 and SampleRate in figure 3.

In the case of CHARM, during the shadowed state the estimated SINR will be that of the previous unobstructed state. Therefore CHARM records a lost packet using the unobstructed state SINR. This will cause the thresholds of that rate to be raised incorrectly when on-line threshold adaptation occurs. When the wireless channel transitions back to the unobstructed state a sub optimal lower transmission rate will be selected. The reason the degradation only gradually starts from 120m onward is due to CHARM's rate fallback mechanism. If no successful transmission is possible due to the shadowed state then CHARM will be stuck at the lowest rate until either a successful transmission or the retransmit limit is reached. Since the optimal rate has been mostly successful prior to the transition, its threshold will not increase because the number of successes far outnumbers the number of failures. When the distance increases the SINR will decrease and fading effects will become more prominent, causing enough failures to trigger a threshold increase at the optimal rate.

For SampleRate during the shadowed state the optimal rate will be disqualified for that averaging window. This means a sub optimal rate will be selected when the channel transitions back to an unobstructed state. Unlike CHARM, the impact starts right away since SampleRate does not take into account transmission successes.




Fig. 4. Hidden terminal node placement, T = transmitter, R = receiver, H = hidden terminal, S = sink for H




Fig. 5. CHARM with threshold adaptation on/off and SampleRate with consecutive losses on/off under hidden terminal interference

C. Hidden Terminal Interference

For this scenario we consider the effects of hidden terminal interference. The widespread popularity of 802.11 networks has led to crowded channels which increases the likelihood of hidden terminal interference. Many rate adaptation algorithms have difficulty with interference because they cannot differentiate packet loss due to collision or normal channel degradation mechanics which leads to sub optimal responses [7].

The setup for this scenario consists of four nodes that are placed in a linear configuration shown in figure 4. We configured the simulation such that a node can only hear neighbors within a 300m range, so the transmitter cannot hear the hidden terminal and vice versa. The hidden terminal transmits at a fixed rate of 11Mbps with a CBR application generating 1, 2, 3.2 and 4 Mbps of interference traffic. We disabled RTS/CTS because previous works have shown RTS/CTS lowers the performance unless under very heavy interference. From figure 5 we can see that much like the EM scenario, there are significant performance degradations for both CHARM and SampleRate.

Interference is similar to EM in that the packet collision is a




Fig. 6. CHARM average throughput with hidden terminal interference but without fading with threshold adaptation on/off

temporary event and the channel conditions prior to and after the event have no relation to the channel conditions during the event. Much like in the case of EM, CHARM will mark a packet lost due to an interference event at the SINR prior to the event causing thresholds to be raised incorrectly. SampleRate will use a lower rate in response to losses which is already known to be a sub optimal response for this scenario.

The original CHARM paper also contains a hidden terminal scenario but did not consider fading effects. We replicate this result by running the same scenario with fading turned off, shown in figure 6. We can see that there is no performance degradation without fading effects. In the original paper the authors claimed that their results were similar even with fading turned on, which we demonstrated to not be the case in more general hidden terminal scenarios. This difference can be explained by the difference between our respective test-bed environments. In the original paper the authors used an emulation test-bed to force a hidden terminal by giving the channel between the transmitter and hidden terminal a very high pass loss. We suspect since they did not consider distances and the channel between the transmitter to the receiver and receiver to the hidden terminal still had high SINR. Since we use a simulation test-bed, our scenario considers the actual distances between nodes. At 170m the SINR between nodes is low enough such that fading effects will have a significant impact on successful packet transmission which is not true if the SINR is higher. These results show the need to fully capture the complexities of the wireless channel when evaluating performance.

V. OPTIMIZATION TECHNIQUES

The main cause of the performance degradations for both CHARM and SampleRate in the previous scenarios is the use of stale channel information. In the case of CHARM it is due to the limitation of requiring a successfully received packet to collect channel measurements. For SampleRate it is the inability to distinguish losses due to interference and shadowing from normal channel degradation. To counteract these problems we propose optimization techniques that will infer when the channel is in a "good state" and treat losses that occur during this state differently.

A. REGS SINR

The *Recovery from Earlier Good State* (REGS) is a cross layer optimization technique originally proposed in a paper by Varshney et al [5] to counteract the memory effect caused by EM. REGS works by saving the protocol state information when the wireless channel transitions to a shadowed state and restoring that information when the channel transitions back. The key assumption is that the states prior and after the shadowed state are similar. Doing this avoids the memory effect by preventing the protocol from using information collected during the shadowed state to make decisions in the unobstructed state. In the original paper it is assumed that REGS can already detect state transitions perfectly and no possible implementation is given. This version of REGS, which

we term REGS Ideal, has already been shown to eliminate the memory effect caused by EM for both SampleRate and TCP.

In this paper we implement a modified version of REGS which we term REGS SINR that uses the novel technique of using estimated receiver SINR to determine state transitions. The reason why transition detection is hard is because when the channel is in the shadowed state packets will be lost completely; therefore no information can be sensed from the wireless channel. REGS SINR works by leveraging this fact, which also implies that *any packet lost when the wireless channel is previously sensed to be in a "good state" condition is due to the channel being in a shadow state*. REGS SINR uses the same receiver SINR estimation method used by CHARM. If a packet becomes lost when the estimated receiver SINR for the channel is above a certain "good state" threshold, that packet is suspected to be lost due to blockage. The "good state" thresholds are set heuristically using the results from figure 1. To determine whether the channel is in a shadowed state we detect if consecutive packets are lost above the "good state" threshold for a duration of at least 20ms, which is the average duration of a fading effect. For the unobstructed state we detect for a successful packet transmitted.

B. IntRec

Interference Recovery (IntRec) is an optimization technique that mitigates the effects of interference on rate adaptation algorithms. The main commonality between EM and interference is they can both be modeled as a binary state system. For interference the "good state" is when the channel does not have cross traffic while the "bad state" is when it does. The channel conditions during the "good state" are distinctly separate and unrelated from the channel conditions during "bad state". During the "bad state" the channel conditions are such that almost no packet can be successfully transmitted. Therefore the transmission rate should only be optimized for the channel during the "good state" while minimizing the harmful effects of the "bad state."

Even though EM and interference have commonalities, due to the random nature of interference REGS SINR cannot be used. The state durations of interference are magnitudes smaller than EM state durations, compare the transmission time of a packet to 200ms. REGS SINR will not be able to properly detect these short term state transitions and thrashing will occur. However what IntRec does is use the same concept of a previously sensed "good state" to do optimizations. What IntRec does is to ignore information collected for packets that are lost above the "good state" threshold.

The main difference between REGS SINR and IntRec is that REGS SINR only restores the previous "correct" information while IntRec ignores "wrong" information. REGS SINR allows the underlying algorithm to fully run during the "bad state" while IntRec does not.

VI. OPTIMIZATION EVALUATION

In this section we implement REGS SINR and IntRec for CHARM and SampleRate and evaluate the resulting perfor-




Fig. 7. CHARM average throughput at various distances with REGS SINR and IntRec under EM using 30, 50, 70 % shadowed state




Fig. 8. SampleRate average throughput at various distances with REGS SINR and IntRec under EM using 30, 50, 70 % shadowed state

mance using the same wireless scenarios described in section 3. For EM the results are shown in figure 7 for CHARM and figure 8 for SampleRate. A performance increase of up to 340% can be seen for CHARM and 220% for SampleRate. We can see that for CHARM both REGS SINR and IntRec work extremely well in eliminating the memory effects caused by EM, while IntRec works better for SampleRate. The reason for this is because REGS SINR granularity is not fine enough to save the state of SampleRate just before four consecutive losses, unlike IntRec which prevents the losses from being recorded in the first place.

The results for hidden terminal interference are shown in figure 9 for both CHARM and SampleRate. In the case of CHARM both IntRec and REGS SINR work well in eliminating the effects of hidden terminal interference with a slight edge to IntRec. There is a performance increase of 293% at 4Mbps of interference which makes an unusable throughput of 34kbps to a more reasonable 147kbps. Similarly to EM, IntRec works better for SampleRate with a performance improvement of 88% at 3.2Mbps of interference.

VII. CONCLUSION

Even though IEEE 802.11-based wireless networks have become common, the rate adaptation component of the standard has yet to be defined. Since a key part of rate adaptation is estimating channel conditions and selecting the proper rate, inaccurate channel estimations will greatly degrade the performance. In this paper, we investigated how complex wireless environments with effects such as EM and hidden




Fig. 9. CHARM and SampleRate average throughput with REGS SINR and IntRec with hidden terminal

terminal interference can cause inaccurate channel estimations for two state of the art rate adaptation algorithms, CHARM and SampleRate. We created high fidelity emulations of CHARM and SampleRate in QualNet and evaluated the impact of these inaccurate estimates for EM and hidden terminal interference scenarios. Our results show that the performances of both CHARM and SampleRate degrade dramatically for both scenarios, mainly due to stale channel information. We present two optimization techniques, REGS SINR and IntRec, to counteract the inaccuracies caused by EM and hidden terminal interference effects. These two techniques increase the performance dramatically for CHARM, up to 340% for EM and 290% for hidden terminal interference. We saw similar improvements of 220% and 88% for SampleRate.

REFERENCES

- [1] J. Bicket. Bit-rate selection in wireless networks. Master's thesis, MIT, 2005.
- [2] <http://www.scalable-networks.com>. Qualnet.
- [3] D. Parsons. *The Mobile Radio Propagation Channel*. John Wiley and Sons Inc., 2001.
- [4] G. Judd, X. Wang, and P. Steenkiste. Efficient channel-aware rate adaptation in dynamic environments. In *Proc. of the ACM MobiSys Conf*, pages 118–131, Breckenridge, CO, June 2008.
- [5] M. Varshney and R. Bagrodia. Performance implication of environmental mobility in wireless networks. In *INFOCOM 2007. 26th IEEE International Conference on Computer Communications*, pages 767–775. IEEE, 6–12 May 2007.
- [6] M. Varshney, J. Zhengrong, M. Takai, and R. Bagrodia. Modeling environmental mobility and its effect on network protocol stack. In *Wireless Communications and Networking Conference, 2006. WCNC 2006*, volume 1, pages 587–592. IEEE, 3–6 April 2006.
- [7] S. Wong, H. Yang, S. Lu, and V. Bharghvan. Robust rate adaptation for 802.11 wireless networks. In *Proc. of ACM MobiCom Conf*, pages 146–157, Los Angeles, CA, Sept 2006.