EXTENDED ABSTRACT: Low-overhead Channel-aware Rate Adaptation

Glenn Judd
Carnegie Mellon University
Pittsburgh, PA, USA
glennj@cs.cmu.edu

Xiaohui Wang
Carnegie Mellon University
Pittsburgh, PA, USA
xiaohuiw@andrew.cmu.edu

Peter Steenkiste
Carnegie Mellon University
Pittsburgh, PA, USA
prs@cs.cmu.edu

ABSTRACT

Current rate selection algorithms are dominated by probe-based approaches that search for the best transmission rate using trial-and-error. When operating over a dynamic channel, probe-based techniques can perform poorly since they inefficiently search for the moving target presented by the constantly changing channel. We have developed a channel-aware rate adaptation algorithm - CHARM - that responds quickly to dynamic channel changes, and significantly outperforms probe-based algorithms in many instances. Unlike previous approaches, CHARM leverages channel reciprocity to obtain channel information without incurring RTS/CTS overhead. Our work shows that channel-aware rate selection is viable, and can significantly outperform probe-based rate adaptation over both static and dynamic channels.

Categories and Subject Descriptors

C.2 [**Computer Systems Organization**]:
Computer-Communication Networks
; C.2.1 [**Computer-Communication Networks**]:
Network Architecture and Design
Wireless communication

General Terms

Measurement,Performance

Keywords

Rate selection

1. INTRODUCTION

When selecting a transmission rate, a wireless device faces a fundamental tradeoff between data-rate and range. Higher transmission rates increase throughput, but reduce the range at which the transmission can be successfully decoded since signal power and channel capacity necessarily decrease with distance. When sending to a particular receiver, the transmitter wants to use the highest transmission rate that can still be decoded with high probability. The problem faced by the transmitter is that two key factors are unknown: the channel transfer function, and the noise level at the receiver. Cellular networks address this problem by incorporating feedback from the receiver to the sender. Wireless data networks provide no standard feedback mechanism, leaving transmitters in the dark. Consequently, most rate selection algorithms to date blindly search for the best possible transmission rate using in-band probing.

We introduce a **ch**annel- **a**ware **r**ate selection algorith**m** - CHARM - that directly leverages signal strength information to achieve good rate adaptation performance in dynamic wireless channels. Since it has access to more information, it adapts more effectively than the currently deployed algorithms that are based on probing. CHARM is also very efficient since it quickly obtains accurate channel information without incurring the additional overhead of RTS/CTS required by earlier proposals.

In this extended abstract, we first describe how a transmitter can estimate path loss and the receiver-side SINR in networks that use commercial network interface cards (NICs). We then describe the CHARM protocol, implementation, and evaluation. We conclude with a discussion of related work and conclusions.

2. ESTIMATING RECEIVER SINR

The probability of a successful packet reception at a receiver is largely determined by the signal-to-interference and noise ratio (SINR) at the receiver. If the transmitter could accurately predict all three factors (signal, noise, interference) at the receiver, it could directly select the best transmit rate without resorting to probing. However, getting the necessary information is complicated because commercial NICs provide only limited information and the receiver-side information is needed by the transmitter. We discuss how we address these challenges in this section.

The received signal strength of a wireless signal in dB can be expressed as [11]:

$$RSS = P_{tx} + G_{tx} - PL + G_{rx} \tag{1}$$

Where RSS is the received signal strength, P_{tx} is the transmit power, G_{tx} and G_{rx} are the transmit and receive antenna gain, and PL is the path loss. P_{tx}, G_{tx} and G_{rx} are properties of the transmit and receive hardware and are generally fixed. Their values can be obtained from the hardware and provided to the transmit side rate selection algorithm,

---

This research was funded in part by NSF under award numbers CCR-0205266 and CNS-0434824.

MobiCom'07, September 9–14, 2007, Montréal, Québec, Canada.

although in practice it is not necessary to know the individual values for G_{tx} and G_{rx}. The only portion of Equation 1 that is difficult to obtain is the path loss PL. It is determined by the signal transfer function between the transmitter and the receiver and it constantly changes as a result of motion that occurs within the environment surrounding the transmitter and receiver. For channel-aware rate adaptation to be viable, we need a method of conveying this information to the transmitter in a low-overhead fashion.

CHARM obtains path loss information by leveraging the Reciprocity Theorem [14], which states: "If the role of the transmitter and receiver are instantaneously interchanged, the signal transfer function between them remains unchanged." In other words, the instantaneous path loss between two nodes is the same in both directions and a transmitter can obtain the path loss to a receiver by measuring the path loss from the receiver to the transmitter. Practically speaking, this means that if the "transmitter" knows the transmit power used by the "receiver", it can estimate the path loss by observing the RSS for packets it receives from the "receiver", taking antenna gains into account (Equation 1). As explained later, CHARM uses a weighted average of past path loss values to estimate the current path loss. For each of these values, we leverage leverage reciprocity to obtain path loss information from the sender to the receiver by simply *passively* observing packets sent in the other direction (e.g. traffic or beacons). Earlier approaches that leverage reciprocity rely on RTS/CTS, effectively *active* probes, to obtain instantaneous path loss measurements.

Wireless cards provide a "Received Signal Strength Indication" (RSSI) for every packet that they receive. In modern wireless cards, there is a (very nearly) linear relationship between RSSI and RSS [4]. So, RSSI is a good approximation of RSS and, combined with transmit power information, it can be used to estimate path loss. In our work, RSSI non-linearity is not an issue since we automatically calibrate SINR thresholds dynamically.

The other two factors in the SINR are noise and interference. Noise is the sum of truly random continuous signals present in the communication band of interest as well as inherent in the receiver itself. Since true noise changes on a large time scale, it can be measured by the receiver and provided to the transmitter. Interference refers to signals present at a receiver that were not generated by the transmitter of interest. Wideband continuous sources of interference can loosely be treated as "noise", but bursty interference is more problematic and nearly impossible to predict. The impact for sources that use a medium access protocol (e.g. based on carrier sense) should be relatively small, but the impact of bursty interference generated by non-compliant sources can be more significant.

Note that RSSI is measured at packet acquisition time, before capture effect can play a role and as a result, the effect of interference on RSSI is generally under 1 dB. The most that we have been able to deliberately affect RSSI using interference is approximately 3 dB. Intuitively, if the interference is stronger than that, the "interfering" packet is received or no reception takes place.

3. CHANNEL-AWARE RATE ADAPTATION ALGORITHM (CHARM)

The CHARM protocol has four components: information gathering, path loss prediction, rate selection, and rate SINR threshold estimation.

Gathering Information - All nodes inform other nodes within their transmission range about the transmit power they use and the noise level they observe. In our implementation, which targets infrastructure 802.11 networks, this is done by introducing an additional 802.11 information element in beacons, probe requests, and probe responses. Nodes also monitor the RSSI for incoming packets (data packets and ACKs) from destination nodes they communicate with. Using the measured RSSI and the transmit power and noise information provided by the sender, they can then estimate the instantaneous path loss (including the fixed antenna gains) for the path from and to the sender using Equation 1.

Predicting Path Loss/SINR - In order to select the most appropriate transmit rate, CHARM tries to predict path loss - and equivalently SINR - at the receiver at the time of each packet transmission using the path loss estimates obtained based on packets received from the intended destination. The traditional approach of predicting future values based on history is to use a moving average of past values as a predictor. However, analysis of traces of path loss collected in a number of environments shows distinct trends on specific timescales, so it is important that the prediction algorithm considers *timing* information. Specifically, more recent samples are more likely to be representative of the current channel conditions than older samples, so they should carry more weight. For this reason, CHARM uses a weighted moving average, where weights are controlled by packet arrival time. We also observed that the RSSI measurements include outliers that have no predictive value, so our algorithm includes a preprocessing phase that filters out the outliers. Trace-based evaluation of the path loss prediction algorithm shows that it is very accurate. The prediction error is mostly under 2 dB.

Rate Selection - Before sending a packet, the sender uses a SINR threshold table (described below) for the intended receiver to determine a set of transmission rates. The Atheros chipset allows the driver to specify several transmission rates which will be used for the initial transmission and for each of the possible retransmissions in the order specified by the driver. This allows the firmware to attempt delivery several times at different rates without contacting the driver for each failure. For the first transmission, the driver picks the highest rate supported for the estimated SINR value, in order to maximize the channel throughput. For retransmissions, lower rates are selected from a fast decreasing rate sequence. There are two reasons for switching to lower rates fairly quickly. First, since the first transmission failed, the first rate may have been too high. Second, we would like to deliver the packet fairly quickly since successful delivery will result in an ACK, which give us more up to date information on the SINR, and minimize the time spent probing the channel. Updated SINR information can then be used to pick a better transmit rate for subsequent packets.

Rate SINR Threshold Estimation - Each transmit rate has an minimum required SINR in order for packet reception to occur with a good probability. Ideally, this value would be the same for all node pairs. However, imperfections in transmit power information, RSSI readings, receiver noise estimation, unreported interference, and multipath effects can affect this threshold. To overcome these issues, CHARM includes a mechanism that automatically calibrates SINR thresholds on-line according to observed performance. The Rate SINR Threshold Estimation module starts with a table that holds default values for the SINR threshold for the different rates. The SINR threshold for each rate is then updated by observing packet success rate as a function of predicted SINR. As these thresholds may vary from receiver to receiver, each transmitter contains a rate SINR threshold set for each receiver that it is communicating with, and updates these thresholds independently.

Implementation - The CHARM rate selection algorithm has been implemented in the Madwifi driver for the Atheros chipset. Besides implementing the functions described in this section, the implementation also deals with issues such as legacy nodes (which do not provide transmit and noise power information).

4. EVALUATION

To demonstrate CHARM's effectiveness we measured its performance against several existing rate selection algorithms in both static and dynamic environments.

4.1 Static Scenarios

We first evaluated CHARM's performance against the three rate selection algorithms provided with the Madwifi driver using a UDP throughput test conducted in a static environment. In this test, the transmitter constantly sent as many UDP packets as possible to the receiver. The receiver recorded the number of packets it received at one second intervals. For most scenarios, we conducted four tests of 20 seconds each. We treat all 80 one second measurements as individual trials and report summary statistics. We repeat this test for four rate selection algorithms: CHARM, AMRR, ONOE, and Sample. We compared rates in 11 different locations located in four buildings in three different geographic locations. The first location - "Home" - is a suburban townhome. The second and third locations - REH and WEH - are university campus buildings with an operational 802.11b/g production network. The fourth location - "Apartment" - is an urban apartment. Note that all the experiments were done "in the wild", so they automatically account for the effects of interference and noise that are naturally present in deployed wireless networks.

Figure 1 shows the results of our tests. In eight locations, CHARM significantly outperforms the best of the other three algorithms. In two locations, CHARM performs essentially the same as the best of the other three algorithms. In one location, the best of the other three algorithms significantly outperforms CHARM. The one location where charm performed poorly was located in a university library. The receiver was located on a metal shelf, and the transmitter was obscured from the receiver. The presense of a large amount of metal shelving resulted in extreme multipath fading.

In general, in poor signal environments, CHARM performs similarly to the best of the other algorithms, though Sample may outperform CHARM's current implementation in a severe multipath environment. In moderate to good signal environments, CHARM significantly outperforms the best of the others.

4.2 Mobile Scenarios

Mobility results in rapid channel variations that are challenging for rate selection algorithms to keep pace with. We compared CHARM against the same three algorithms (AMRR, ONOE, and Sample) in two mobile scenarios. In each scenario, the receiver was stationary while the transmitter moved within range of the receiver for 40 seconds. Each algorithm was tested two times for each scenario. Figures 2 and 3 show the results. CHARM significantly outperforms all other algorithms for the vast majority of the time. There is one small region where Sample outperforms all others by remaining aggressive when the channel degrades, but this is short-lived. For almost the entire trace, CHARM's ability to quickly gain an accurate picture of channel state translates into dramatically better performance.

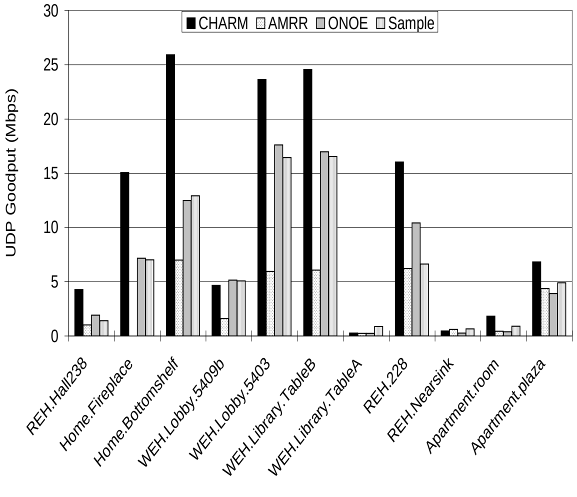

Figure 1: Median Throughput

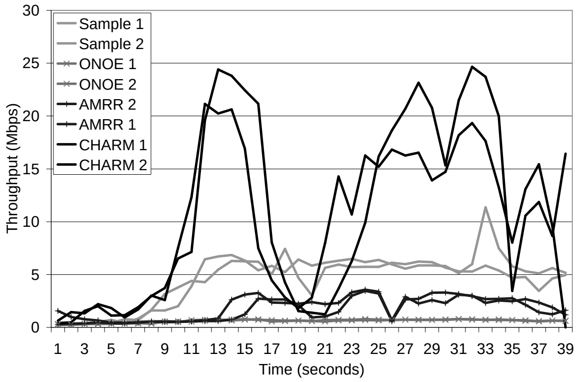

Figure 2: Mobile Throughput Scenario 1

5. RELATED WORK

The problem of transmission rate selection has been widely studied. We broadly categorized these approaches into probe-based, SINR-based, and hybrid, although hybrid elements are present in many probe and SINR-based algorithms.

Probe-based rate selection algorithms use information about successful or failed packet reception as implicit indicators of reception conditions at the receiver [15, 7, 8, 1, 13, 9]. These algorithms typically use in-band probing via user data packets. 802.11 acks provide the transmitter with

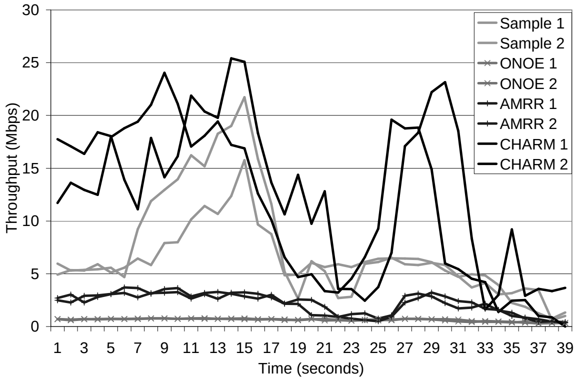

Figure 3: Mobile Throughput Scenario 2

knowledge that reception occurred; ack-timeouts are taken as an indication that reception did not occur, though this may not be the case if it is the ack packet that is lost. The advantage of probe-driven approaches is simplicity, and the ability to implicitly take into account complex factors affecting reception. A key disadvantage is the speed at which channel information can be obtained. From the transmitter's perspective, each transmission attempt results in either a success, or a perceived failure. Another major disadvantage of probe-based rate adaptation is the inability to distinguish between the causes of perceived transmission failure; all the transmitter knows is that it did not correctly receive an ACK. For example, packet losses caused by collisions generally do not justify reducing the transmission rate.

In contrast to probe-based approaches, **SINR-based** approaches use signal metrics provided by the wireless devices to select the transmission rate [3, 12, 10]. The algorithms typically rely on the RTS/CTS mechanism to provide instantaneous receiver-side SINR information to the transmitter. In theory, knowing the SINR at the receiver would allow the transmitter to directly set the transmission rate without wasting precious time probing. However, the use of the RTS/CTS mechanism to communicate the receiver SINR to the transmitter introduces significant overhead, which CHARM avoids. Moreover, relying on a single unfiltered SINR measurement can potentially result in poor rate selection. Channel reciprocity has also been exploited in a number of other context, including cellular TDD systems [5] and the 802.11n [6] standard for MIMO.

[2] introduces a hybrid approach. It is effectively a probe-based technique in which SINR information is used to restrict the set of rates that can be used by the probing algorithm. This approach shares some elements with CHARM, but there are significant differences. CHARM uses SINR information as the primary source of information for rate selection. Historical information on success or failure of packet transmission is used indirectly to optimize the threshold used by the SINR-based algorithm.

6. CONCLUSION

We have developed a channel-aware rate adaptation algorithm that quickly obtains accurate channel state information, and, unlike earlier channel-aware efforts, leverages channel reciprocity to eliminate the need for RTS/CTS exchanges. We use time-aware signal prediction technique to predict current channel information based on past observations while avoiding the pitfall of using stale channel information. In addition, we have developed techniques for automatically calibrating SINR thresholds, as well as methods for sharing remote transmit power and noise information. Experiments conducted in both relatively static and mobile scenarios show that CHARM performs very well compared with commonly used rate selection algorithms.

7. REFERENCES

[1] J. C. Bicket. Bit-rate selection in wireless networks, February 2005.

[2] I. Haratcherev, K. Langendoen, R. Lagendijk, and H. Sips. Hybrid rate control for ieee 802.11. In *Proc. of MobiWac 2004. Philadelphia, PA*. ACM, October 2004.

[3] G. Holland, N. Vaidya, and P. Bahl. A Rate-Adaptive MAC Protocol for Multi-hop Wireless Networks. In *Proc. of MobiCom2001. Rome, Italy*, September 2001.

[4] G. Judd and P. Steenkiste. A simple mechanism for capturing and replaying wireless channels. In *E-Wind 2005*, Philadelphia, PA, August 2005. ACM.

[5] V. Jungnickel, U. Kruger, G. Istoc, T. Haustein, and C. Helmolt. A mimo system with reciprocal transceivers for the time-division duplex mode. In *AP-S International Symposium 2004*, June 2004.

[6] M. Kuhn, A. Ettefagh, I. Hammerstrom, and A. Wittneben. Two-way communication for ieee 802.11n wlans using decode and forward relays. In *ACSSC '06*, Pacific Grove, CA, USA, August 2006.

[7] T. T. M. Lacage, M. H. Manshaei. Ieee 802.11 rate adaptation: A practical approach. In *Proc. of MSWiM 2004. Venice, Italy*, pages 126–134. ACM, October 2004.

[8] Madwifi. Multiband Atheros Driver for WiFi.

[9] D. Qiao and S. Choi. Fast-responsive link adaptation for ieee 802.11 wlans. In *Proc. of ICC 2005. Seoul, Korea*. IEEE, May 2005.

[10] S. C. L. Qixiang Pang, Victor C.M. Leung. A rate adaptation algorithm for ieee 802.11 wlans based on mac-layer loss differentiation. In *2nd International Conference on Broadband Networks*, pages 659–667. IEEE, October 2005.

[11] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice-Hall, Englewood Cliffs, NJ, 2002.

[12] B. Sadeghi, V. Kanodia, A. Sabharwal, , and E. Knightly. Opportunistic media access for multirate ad hoc networks. In *Proc. of MobiCom 2002. Atalanta Georgia, USA*. ACM, September 2002.

[13] S. L. Starsky H.Y. Wong1 Hao Yang and V. Bharghavan. Robust rate adaptation for 802.11 wireless networks. In *Proc. of the MobiCom 2006. Los Angeles, CA*. ACM, September 2006.

[14] C. Tai. Complementary reciprocity theorems in electromagnetic theory. *IEEE Trans. on Antennas and Propagation*, 40(6):675–681, 1992.

[15] V. van der Vegt. Auto Rate Fallback.