

Overcoming MAC Overhead Using Packet-Size Dependent Channel Widths[†]

Vijay Raman, Fan Wu, Brian Proulx, and Nitin H. Vaidya
University of Illinois at Urbana-Champaign
{vraman3, fwu24, proulx1, nhv}@illinois.edu

I. INTRODUCTION

The present day communication networks have to carry traffic with different packet sizes. For instance, VoIP traffic and TCP ACKs (that are generated by web browsing and HTML traffic) are short packets of few 100 bytes, while file transfers and other TCP sessions may involve packets of the order of 1000 bytes. An internet traffic analysis by CAIDA in 2008 shows that over 55% of the traffic in the internet are smaller than 100 bytes and this trend has not changed over the past ten years [1]. This suggests that a variety of packet size mix will exist in the future communication networks.

Most of the present day wireless communication techniques follow a random access scheme, where the channel is first assessed to be free before a packet transmission to avoid collisions (e.g., DIFS in 802.11). If the channel is not free, the nodes backoff until the channel becomes free again. The time spent on clear channel assessment and backoff is independent of the packet size or the rate of transmission. Hence, these overheads associated with a packet transmission are known as rate-independent or bandwidth independent overheads. If, for instance, P_l (in bits) denotes the packet payload size, T (in seconds) denotes the channel time consumed by the rate-independent overhead associated with each transmission, and R (in bits per second) denotes the transmission rate, then adopting a simplified model, $\frac{TR}{P_l+TR}$ fraction of channel capacity is wasted as the rate-independent overhead [2]. Observe that the channel wastage is higher when the packet payload size is small or when we use higher rates of transmission. This suggests that we use a lower rate of transmission for shorter packets than compared to a large packet. However, simply reducing the rate of transmission may result in the short packets occupying the channel for a longer time, which may be unfair for any long packets that are to be transmitted. Furthermore, using a lower rate of transmission wastes the channel capacity.

In this work, we propose to partition the channel into a narrow channel and a wide channel. The narrow channel is used for transmitting the short packets and the wide channel is used for transmitting the longer packets. We intend to use multiple radios, one each for the different channel partitions. Narrow width channels have a reduced capacity, which as

a result lowers the maximum transmission rate achievable on these channels. As a result, the channel wastage in rate-independent overhead can be reduced. However, it is not straightforward as to how much bandwidth to allocate for short packet transmissions. This is because, if a node predominantly transmits long packets with very little short packets, then the capacity lost in partitioning the channel may cause a negative effect on the throughput achieved by the long packets. On the other hand, if the node generates more short packets than long packets, then the bandwidth allocated for the short packets may not be sufficient, which may result in buffer overflow and packet loss at the sender side. It is therefore important to determine the bandwidth required for each of the packet sizes. Furthermore, it is also important to understand when to partition the channel, depending on the amount of short and long packets generated in the network.

Some of our initial analysis has provided good intuitions for developing protocols for selecting the channel widths based on the packet sizes. Based on these intuitions, we develop a centralized protocol for determining the channel widths, assuming a single hop infrastructure mode network. In this extended abstract, we provide an overview of the protocol. We also discuss the means for extending our protocol to a multihop, ad-hoc setting and some related issues. We intend to implement the protocols on a real testbed to show the feasibility of our protocols using commodity hardware.

II. PROBLEM MOTIVATION

In this section, we demonstrate the benefit of choosing variable-width channels based on packet sizes. First, we wish to understand the amount of capacity loss when higher rates are used for short packets. For this, we generate packets of various sizes ranging from 100 bytes to 1500 bytes and plot the capacity loss calculated at various fractions of bandwidths. If α is the fraction of bandwidth allocated for the packet transmission and $DIFS, SIFS$ represent the inter-frame spacing in IEEE 802.11 (chosen to be $34 \mu s$ and $16 \mu s$ respectively, considering a slot duration of $9 \mu s$), the capacity loss, C_{loss} is calculated using the following formula,

$$C_{loss} = \frac{(DIFS + SIFS) * \alpha R}{P_l + (DIFS + SIFS) * \alpha R}$$

In this equation, R is the maximum rate of transmission, which at a bandwidth of 20 MHz (802.11 channel width) is 54 Mbps. We assume that the rate of a packet transmitted

[†]This work was supported in part by NSF grant 06-27074 and in part by US Army Research Office grant W911NF-05-1-0246. Any opinions, findings, and conclusions expressed here are those of the authors and do not necessarily reflect the views of the funding agencies or the U.S. government.

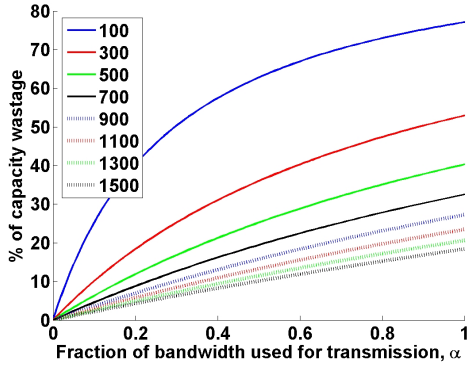


Fig. 1. Percentage of capacity loss as a function of fraction of bandwidth used.

at α fraction of the bandwidth is also scaled by α . The C_{loss} values for the different packet sizes are shown in Figure 1. We observe from the plot that shorter packets experience higher capacity loss when they are transmitted at higher fractions of bandwidth than longer packets. In particular, we observe that for a 100 byte packet transmitted at the full bandwidth ($\alpha = 1$), the capacity loss is above 80%, whereas it is lower than 20% for a 1500 byte packet. We also observe that shorter packets experience fewer capacity loss when they are sent at narrower bandwidths. This suggests that choosing bandwidth based on packet sizes can lower capacity loss.

III. NETWORK MODEL

We assume a single hop managed-mode network consisting of a set of static wireless nodes controlled by an access point (AP). For simplicity, we only consider the uplink communication between the wireless nodes and the AP in our discussions, however our model is generic and can also be applied to a downlink communication. We assume that the available spectrum can be split into multiple sub-channels, each of varying widths. Furthermore, we assume that the nodes and the AP are equipped with multiple radios. The wireless radios in a node are capable of transmitting over any one of the sub-channels at any instant of time, and are capable of switching across sub-channels. We assume that the sub-channels have sufficient guard band between them, which does not necessitate the transmissions on the different sub-channels to carrier sense each other. Figure 2 shows an example where a 20 MHz channel is split in two possible ways, (a) two 10 MHz channels, and (b) a 5 MHz and a 15 MHz channel. Note that the center frequencies of the sub-channels change depending on their widths.

We assume that the system may involve several different traffic types generating packets of different sizes. We categorize these traffic into multiple classes based on their packet sizes. The channel widths for each of the classes are chosen proportional to the effective arrival rates of the packets.




Fig. 2. An example where a 20 MHz channel is split into (a) two 10 MHz channels, and (b) a 5 MHz and a 15 MHz channel.

IV. CENTRALIZED ALGORITHM

In the centralized case, the bandwidth partition values are decided by the AP. Note that for successful communication between a pair of nodes, they have to be communicating on the same sub-channel. Thus, the bandwidth partitions chosen by an AP has to be used by all the N clients controlled by this AP. The AP chooses the bandwidth partitions from a set of l bandwidth pairs, $\{(BW_{11}, BW_{12}), (BW_{21}, BW_{22}), \dots, (BW_{l1}, BW_{l2})\}$, such that $BW_{i1} + BW_{i2} \leq BW$, where BW is the bandwidth of the channel being partitioned. The wireless radios in the clients and the AP can operate either on any of the bandwidths chosen from this set or the full bandwidth BW of the channel.

The following specifications are used by our algorithm:

1. Every client i estimates the arrival rate of packets at its side, λ_i and computes the percentage of short packets β_i using the formula, $\beta_i = \frac{\text{No. of packets of size} \leq P_{th}}{\text{Total no. of packets}}$, where P_{th} is a packet size threshold, such that packets smaller than P_{th} are considered short and are otherwise considered long. The clients then send the estimate of arrival rate and the percentage of short packets to the AP periodically every T_{cl} seconds.

2. The AP calculates the aggregate percentage of short packets from all the clients in the network, using $\beta = \frac{\sum_i \beta_i \lambda_i}{\sum_i \lambda_i}$.

3. The AP then chooses the bandwidth allocated to the short packets and sends $[(BW_{j1}, BW_{j2})]$ to all the clients, where $\frac{BW_{j1}}{BW}$ is a bandwidth value that closely approximates β .

4. Once the bandwidth partitions are received by the clients, they use the bandwidth BW_{j1} for short packets (i.e. $< P_{th}$) and the bandwidth BW_{j2} for long packets.

5. To maintain the signal-to-noise ratio (SNR) at the receiver constant, the clients scale their transmission power according to the bandwidth used for transmission. For the same reason, the carrier sense threshold, used by the nodes to assess if the channel is clear before a transmission, is also scaled accordingly.

Note that the AP estimates the bandwidth partitions based on the packets received from all the clients and all the flows. To enable new clients that may later join the AP's network, the AP sends the beacon packets (or neighbor advertisement packets) on the full bandwidth BW along with the information on the current bandwidth partition. Thus the new client can start using the new partitions right away. The AP can then re-calculate the bandwidth partitions for the whole network based on the packets that the new client generates. If any of the channel queue is full at a client, then the client can choose to send any additional packets arriving at that channel queue through the other channel.

A. Extensions to a Distributed Scenario

Our centralized algorithm can be easily modified to a distributed setting suitable for a multihop network. In the distributed case, the bandwidth partitions are estimated by the receive nodes, which for instance, may be the next hop node for a flow. However, the nodes that send the packets to a common next-hop node j , have to send their estimate of the arrival and transmission rates of the packets of only those flows that are sent through j . Thus, different flows in a node can be assigned different bandwidth values depending on the next-hop of a flow. The nodes, therefore, may have to switch across different bandwidth pairs for transmitting the packets belonging to the different flows. Furthermore, every hop of a given flow may be using a different bandwidth values. Note that a single flow targeted at a given next-hop node have to use two different bandwidths, one for the short packets in the flow and the other for the long packets in the flow. Thus, two radios are required for simultaneously transmitting the packets of a flow on the two bandwidths. Furthermore, two more radios are required per node for receiving packets sent by other nodes on two bandwidths. Thus, our distributed algorithm requires that every node is equipped with at least four radios. When more than two partitions are used per flow, more radios may be required.

Few interesting problems arise in the case of a distributed setting as summarized below:

a) Carrier sensing across different bandwidths: The nodes in the case of a distributed setting may be using different bandwidths on the same channel spectrum. An important aspect, therefore, that need to be considered in a distributed setting is the means for carrier sensing transmissions on all the possible bandwidth pairs. One straightforward heuristic will be to carrier sense every possible bandwidth pair before initiating a transmission. The carrier sense thresholds, of course, have

to be scaled according to the bandwidth, as mentioned in the centralized case. This mechanism, however, can be expensive due to the associated latencies. We wish to explore effective alternatives to this simple approach.

b) Interference-aware bandwidth selection: We illustrate this problem using the example in Figure 3. The figure (a) shows two transmissions on the same frequency spectrum, one from node A to B, and the other from node D to C. The transmissions, however use only the bandwidths that




Fig. 3. An example to illustrate interference-aware bandwidth selection. The portion of bandwidth used by the transmissions is shaded black. A cross mark indicates that the transmissions cannot take place simultaneously and a check mark indicates that the two transmissions can be scheduled simultaneously.

are shaded black. These two transmissions cannot take place simultaneously (indicated by a cross mark), as otherwise they will interfere with each other since their bandwidths overlap. If however, the bandwidth are chosen as in (b), then the two transmissions can be scheduled simultaneously without interfering with each other, as their bandwidths do not overlap. Thus, (b) can achieve a higher system throughput than (a). We wish to explore more on efficient bandwidth selection algorithms that are interference-aware.

V. CONCLUSION

In this work, we have proposed to partition a channel into a narrow and a wide subchannel for overcoming MAC overheads. The narrow subchannel is used for sending short packets and the wide channel is used for sending long packets. We have proposed a centralized algorithm for partitioning the channel and have discussed some interesting problems when extending this to a distributed algorithm.

REFERENCES

- [1] (2009) Packet size distribution comparison between internet links in 1998 and 2008. CAIDA Research. [Online]. Available: http://www.caida.org/research/traffic-analysis/pkt_size_distribution/graphs.xml.
- [2] X. Yang and N. Vaidya. (2006, March) Spatial backoff contention resolution for wireless networks. [Online]. Available: <http://www.crhc.uiuc.edu/wireless/groupPubs.html>.