2010

On maximizing VoIP capacity and energy conservation in multi-rate WLANs

Kwan-Wu Chin
University of Wollongong, kwanwu@uow.edu.au

Abstract
This letter highlights a key problem that arises when frame aggregation, a popular method for boosting VoIP capacity, is used in Wireless Local Area Networks (WLANs) that employ asynchronous power save mode (PSM). Specifically, it shows how the PSM proposed by the Wi-Fi Alliance renders frame aggregation ineffective. It then proposes a novel opportunistic scheduler that restores the benefits of frame aggregation whilst ensuring stations have minimal energy expenditure.

Keywords
era2015

Disciplines
Physical Sciences and Mathematics

Publication Details
K. Chin, "On maximizing VoIP capacity and energy conservation in multi-rate WLANs," IEEE Communications Letters, vol. 14, (7) pp. 611-613, 2010.

On Maximizing VoIP Capacity and Energy Conservation in Multi-Rate WLANs

Kwan-Wu Chin

Abstract—**This letter highlights a key problem that arises when frame aggregation, a popular method for boosting VoIP capacity, is used in Wireless Local Area Networks (WLANs) that employ asynchronous power save mode (PSM). Specifically, it shows how the PSM proposed by the Wi-Fi Alliance renders frame aggregation ineffective. It then proposes a novel opportunistic scheduler that restores the benefits of frame aggregation whilst ensuring stations have minimal energy expenditure.**

***Index Terms*—VoIP capacity, IEEE 802.11, WLANs, opportunistic scheduling.**

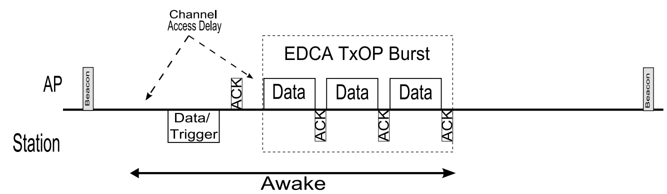

Fig. 1. Wi-Fi multimedia (WMM) power save (PS) extension being used in IEEE 802.11e WLANs, where enhanced distributed channel access (EDCA) is used for channel contention.

I. Introduction

The Voice over IP (VoIP) capacity of Wireless Local Area Network (WLANs) is of key concern to enterprises that plan to converge their data and voice networks. Unfortunately, in practice, the VoIP capacity of WLANs is very low. For example, Wang et al. [1] showed that IEEE 802.11b and 802.11g WLANs are only able to support 12 and 60 VoIP sessions respectively. The main reason is the high overheads associated with the transmission of small VoIP packets. In fact, higher layer signaling amount to an overhead of 680% in IEEE 802.11b WLANs; in the best case, these overheads reduce to 200% when using IEEE 802.11g, but remain at 400% in most cases. Besides that, the constant bit rate nature of VoIP sessions exacerbates collisions, which in turn causes the Rate Adaptation Algorithm (RAA) at the sender to transmit at a lower data rate. As a solution, the authors of [1] propose to combine n small VoIP packets and multicasts only one aggregated frame to stations, which has the affect of increasing VoIP capacity by 100%. This solution, however, considers all stations to be awake at all times.

Energy conservation is a critical issue to consumers with Wi-Fi enabled, power constrained, portable devices. For example, the talk time of Apple's iPhone reduces from 14 hours to eight hours when both the cellular and WLAN interface are switched on [2]. The is because its wireless Network Interface Card (NIC) consumes 280 mA and 204 mA during transmission and reception respectively [3]. However, it only consumes 14 mA when sleeping. Therefore, to save energy, its wireless NIC needs be off for as long as possible. This, however, is challenging as VoIP traffic have packet generation intervals in the tens of milliseconds. Fortunately, they can tolerate small losses and delays.

In 2005, the Wi-Fi Alliance proposed a power saving mode extension that has better supports for real-time data whereby stations are allowed to retrieve their packets from the Access Point (AP) asynchronously, as opposed to periodically as per its AP's beacon period. Figure 1 shows the operation of Wi-Fi Multmedia (WMM)'s Power Save (PS) extension [4]. Upon waking up, a station sends a trigger frame, which could be piggybacked on a data frame. The AP then sends an acknowledgment packet from the AP if the station's trigger frame is received successfully. After that, the AP begins transmitting the station's buffered data packet. Notice that data packets are sent in a burst, which has the advantage of removing delays associated with channel access. Once a data packet with the "more" flag unset is received, the station goes to sleep.

A critical problem that arises when using WMM-PS, and similar PSMs, is that they negate the performance benefits of frame aggregation as stations are required to be awake at all times to receive aggregated packets. Hence, APs are only able to aggregate packets on a per station basis, which means all downlink transmissions will again incur the high cost associated with transmitting small VoIP packets [1]. To this end, in the next section, we propose an algorithm that restores the benefits of packet aggregation. Note that, this problem does not occur if stations wake-up synchronously, as is the case with legacy PSM. However, as pointed out in [4], legacy PSM is not suitable for VoIP data. More importantly, it does not allow fine grain control of stations sleep times, which is needed to cope with the real-time requirement of voice calls and varying delay budgets; to be explained later.

II. Opportunistic Scheduler

Firstly, for each VoIP call from station A, its delay budget (D_A) is defined as the remaining time before the next packet is to be played or before the call experiences a discernible drop in voice quality. For example, if we assume the network delay to be 100ms and a tolerable delay of 150ms, the station making the call would then inform its AP that its link budget is 50ms.

Secondly, we define eight groups, $G_{i=1...8}$, that correspond to the data rates of IEEE 802.11a. It is also possible to define

Manuscript received March 24, 2010. The associate editor coordinating the review of this letter and approving it for publication was G. Lazarou.

K.-W. Chin is with the University of Wollongong, NSW, Australia (e-mail: kwanwu@uow.edu.au).

Digital Object Identifier 10.1109/LCOMM.2010.07.100468

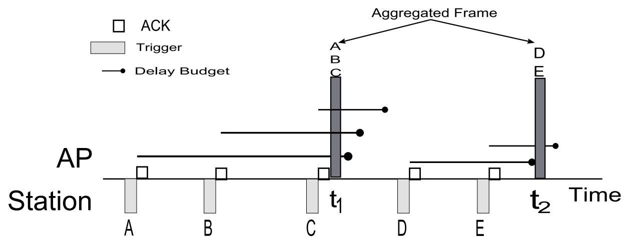

Fig. 2. Opportunistic scheduling.

groups based on different data rate ranges. For example, an AP may have five groups corresponding to stations with data rate less than 10 Mb/s, between 10 and 20 Mb/s, and so forth. The data rate of each group, or the lowest data rate in the range, is then used by the RAA as the initial transmission rate. If the AP fails to solicit an acknowledgment, it attempts the next lower data rate until it reaches the maximum retry count.

In a nutshell, the proposed scheduler, see Algorithm 1, exploits the delay budget of each station to maximally aggregate packets, and thus prolongs the lifetime of stations, and in turn allow a WLAN to have a higher VoIP capacity. Figure 2 illustrates how our algorithm forms an aggregated frame. Without loss of generality, we assume station A, B and C have the same data rate. Initially, all groups do not have any aggregation time. Hence, stations send a trigger whenever they expect the next VoIP packet to arrive at the AP; e.g., every 20ms. Upon receiving a trigger, say from station A, the AP records the station's delay budget and its data rate. The algorithm then searches for a group that matches Station A's data rate, say G_1. Given that group G_1 has no aggregation time, Station A is informed to wake up at time $t_1 = t + (D_A - \nu)$, where t is the current time and ν is the time to transmit an aggregated packet at the base rate. Assume that the trigger from station B is the next to arrive. The algorithm then searches for an aggregation time in G_1 that expires closest to $t + (D_B - \nu)$, and also has sufficient room to store station B's packet; e.g., assuming the G.711 codec, an IEEE 802.11 payload is able to hold approximately 14 packets of 160 bytes. In this case, that time happens to be t_1. The algorithm is then repeated for the trigger from station C.

Note that the aggregated frame, which contains three packets, is constructed similarly to [5]. That is, the AP prepends an aggregation header describing the set of packets and their corresponding packet length. Also, the aggregated frame is unicast to one of the stations randomly, and hence solicits only one acknowledgment (ACK) frame. Moreover, this ensures the aggregated packet is not transmitted at the base rate, which is the case if it was a multicast packet.

III. Simulation Methodology

We used *ns-2* (v2.33) [6] to validate our scheduler. We also employed the new IEEE 802.11 MAC and physical layer extensions implemented by Chen et al. [7], which model the noise floor, physical layer header, transmission and processing of preamble of each packet, and capture effects accurately. In our simulation, all stations are placed uniformly within a

Algorithm 1: Opportunistic scheduler.

`/* Delay budget, station address and data rate. */`

input : D_{sta}, A_{sta}, R_{sta}
output: A_{sta}'s wake up time.
G = FindGroup(R_{sta})
begin
  `/* AggrTxTime is an associative array. */`
  if *G.AggrTxTime is empty* **then**
    G.AggrTxTime$[t + (D_{sta} - \nu)] \cup A_{sta}$
  else
    `/* Get a time closest to sta's delay budget, and is not full. */`
    BigK = FindKey(G.AggrTxTime, $t + (D_{sta} - \nu)$)
    if *BigK* **then**
      G.AggrTxTime[BigK] $\cup A_{sta}$
    else
      G.AggrTxTime$[t + (D_{sta} - \nu)] \cup A_{sta}$
end

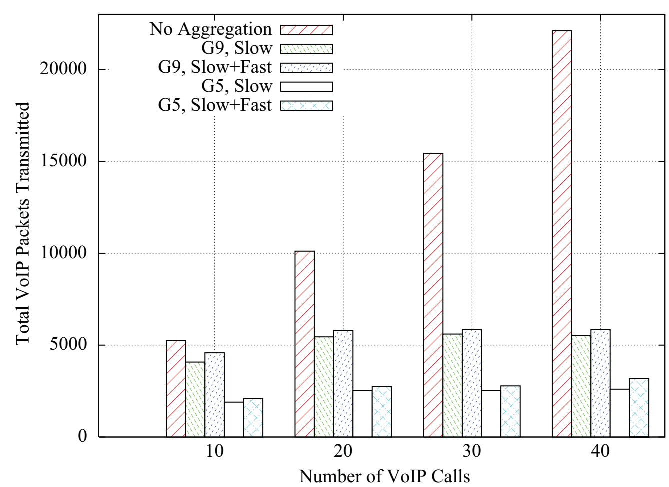

Fig. 3. Total number of packets transmitted.

distance of 250 meter around an AP. Each station determines its transmission rate according to the 2-ray ground model; viz. slow fading. To investigate the impact of channel variation, we vary the signal strength to each station randomly by $\pm$ 20 dBm; viz. fast fading. A station discards a packet after attempting to transmit a packet seven times. Another source of packet loss is when packets missed their playout time, which is set to 150 millisecond in all our experiments. Each station emits one VoIP call. Hence, the number of stations equates to the number of VoIP calls. Each VoIP call consists of two Constant Bit Rate (CBR) flows; one that originates from the station, and the other from the AP, which simulates the corresponding peer of each VoIP call. Both CBR flows start randomly, and emit a 160 bytes packet every 20ms; as per the G.711 codec. In addition, we randomly delay the packets originating from the AP by 50 to 130 millisecond to simulate delays incurred by packets when they traverse the Internet. Hence, the delay budget of stations vary between 100 to 20 milliseconds.

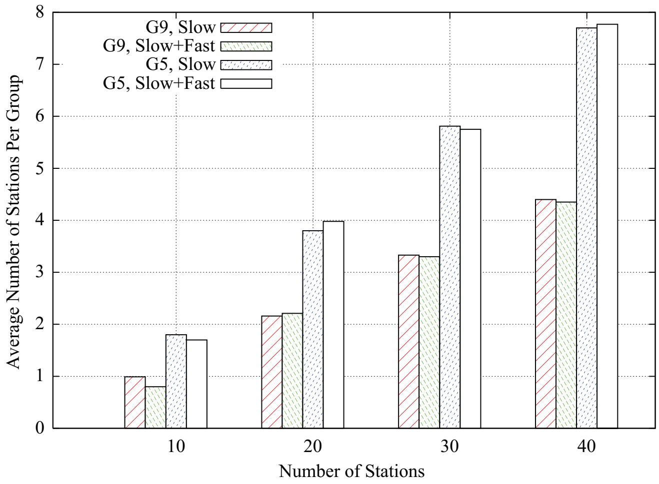

Fig. 4. Average number of stations in each group.

IV. RESULTS

Figure 3 shows the average number of VoIP packets transmitted over a 10 seconds period. In all cases, our scheduler results in fewer packet transmissions. Specifically, when there are only 10 stations, there are 23% to 64% fewer packets on the channel. These savings increase up to 88% when there are 40 stations. In cases with slow and fast fading, when there are nine groups, stations tend to transmit 10% more packets due to channel errors, which require multiple retransmissions. However, if there are only five groups (as described in Section II), this value drops to approximately 6% as lower data rates are used, which have higher coding gains. This, however, increases packet transmission times, but result in fewer retransmissions. Apart from that, in both cases, attempting transmission at the group rate first yields 30% to 50% reduction in air time as compared to multicasting at the base rate, as per the IEEE 802.11 specification.

Figure 4 shows the average number of stations in each group over a 10 second period. As expected, there are more stations when the group size is five. Notice that channel condition does not have a significant impact on group membership. In both cases, performance improves significantly with increasing number of stations.

V. CONCLUSION

This proposed solution restores the use of frame aggregation, and hence improves VoIP capacity in WLANs using asynchronous PSM. In addition, it improves performance by transmitting aggregated packets at their group rate, which has the effect of reducing air-time. Lastly, it is not specific to WMM-PS and can be applied to other asynchronous PSMs with little or no modification.

REFERENCES

[1] W. Wang, S. C. Liew, and V. O. Li, "Solutions to performance problems in VoIP over a 802.11 wireless LAN," *IEEE Trans. Veh. Technol.*, vol. 54, pp. 366–376, Jan. 2005.

[2] V. Namboodiri and L. Gao, "Towards energy efficient VoIP over wireless LANs," in *ACM MobiHoc*, Hong Kong, May 2008.

[3] L. M. Feeney and M. Nilsson, "Investigating the energy consumption of a wireless network interface in an ad hoc networking environment," in *IEEE Infocom*, Alaska, USA, May 2001.

[4] WiFi Alliance, "WMM power save for mobile and portable Wi-Fi certified devices," Dec. 2005.

[5] P. Verkaik, Y. Agarwal, R. Gupta, and A. C. Snoeren, "Softspeak: making VoIP play well in existing 802.11 deployments," in *USENIX NSDI*, Boston, USA, Apr. 2009.

[6] "The Network Simulator NS-2," http://www.isi.edu/nsnam/ns/.

[7] Q. Chen, F. Schmidt-Eisenlohr, D. Jiang, M. Torrent-Moreno, L. Delgrossi, and H. Hartenstein, "Overhaul of IEEE modeling and simulation in ns-2," in *ACM MSWiM*, Crete Island, Greece, Oct. 2007.