

A Practical Unacknowledged Unicast Transmission in IEEE 802.11 Networks

Hyun Yang, Jin Seok Yun, Jun Seok Oh and Chang Yun Park

School of Computer Science and Engineering, Chung-Ang University

221 Heuk Seok-Dong, Dongjak-Gu, Seoul, Korea

[e-mail : {yanghyun, yunjs, ohjunseok}@cnlab.cse.cau.ac.kr, cypark@cau.ac.kr]

*Corresponding author: Chang Yun Park

*Received August 22, 2010; revised October 11, 2010; revised February 16, 2011;
accepted February 22, 2011; published March 31, 2011*

Abstract

In current IEEE 802.11 wireless LAN, every unicast transmission requires an ACK from the receiver for reliability, though it consumes energy and bandwidth. There have been studies to remove or reduce ACK overhead, especially for energy efficiency. However none of them are practically used now. This paper introduces a noble method of selective unacknowledged transmission, where skipping an ACK is dynamically decided frame by frame. Utilizing the fact that a multicast frame is transmitted without accompanying an ACK in 802.11, the basic unacknowledged transmission is achieved simply by transforming the destination address of a frame to a multicast address. Since removing ACK is inherently more efficient but less strict, its practical profit is dependent on traffic characteristics of a frame as well as network error conditions. To figure out the selective conditions, energy and performance implications of unacknowledged transmission have been explored. Extensive experiments show that energy consumption is almost always reduced, but performance may be dropped especially when TCP exchanges long data with a long distance node through a poor wireless link. An experiment with a well-known traffic model shows that selective unacknowledged transmission gives energy saving with comparable performance.

Keywords: IEEE 802.11, wireless LAN, energy-efficient communication, unacknowledged transmission

1. Introduction

In most wireless networks, a unicast transmission requires an ACK (Acknowledgement) message from the receiver. Although it consumes energy and bandwidth, this ACK has less meaning in some circumstances, for example, when transmission error hardly occurs. In order to reduce this resource waste, some wireless protocols such as 802.15.4 provide an option, so-called *unacknowledged transmission mode* [1], but its usage has been rarely studied. Most of all, an ACK frame should be always transmitted after every unicast frame in current 802.11, the most popular wireless network.

Our study is motivated from the following observations. First, 802.11 is so widely used that almost every home or office room has its own AP. In such an environment only a few nodes access the AP within a few meters, and thus frame errors are rare. Second, some applications such as sensor and multimedia applications have relaxed reliability requirements, so that some data are allowed to be missed. Third, contrary to unicasting, 802.11 multicasting does not accompany an ACK, and hence there exists a room not to involve an ACK in transmission. Therefore, it might be feasible and practically effective to make the current ACK-based 802.11 unicasting *less strict but more efficient*.

The research theme is to save resources, especially, *energy* by removing ACKs in unicast communication in 802.11 LAN. The proposed scheme, called simply *unacknowledged unicasting*, should be applicable to the existing hardware and software, transparently. To apply unacknowledged transmission only when it is expected practically beneficial, the selection rules are designed based on the condition of a link and the traffic parameters of a frame.

This paper introduces a simple method to remove an ACK; it transfers a unicast traffic of the higher layers at 802.11 not by a unicast frame but by a multicast frame. We have developed a mechanical 1-to-1 transformation between a unicast MAC address and a unique multicast address. Just before the transmission of a unicast frame, its destination MAC address is transformed to the corresponding multicast address. On the reception of the frame, the destination address is transformed back to the original MAC unicast address and then processed as usual; no ACK frame is transmitted. The proposed scheme could be implemented easily by adding some codes to the LAN driver. In the paper, all explanations are made based on 802.11 infrastructure networks, and they would be the same as well in 802.11 ad hoc networks.

The energy and performances implications of unacknowledged transmission have been explored through extensive measurements and simulations. For multimedia applications using UDP, unacknowledged unicasting could provide energy saving without noticeable quality degradation. For reliable applications using TCP, the effectiveness varies depending on network and traffic conditions. With respect to energy consumption for communication, energy saving could be achieved in most cases. The throughput is also improved in case of intranet communication through a high quality wireless link. This evaluation could give some general insights on how to utilize unacknowledged transmission.

The expected side-effects are performance degradation caused by end-to-end error and congestion controls due to the lack of immediate link-level action, i.e., ACK. Considering the current and future dissemination of 802.11, we believe that our scheme gives more chances of gain than harm. Moreover, those harmful cases are predictable and hence avoidable by switching back to the original 802.11. Therefore, "*selective unacknowledged unicast*

transmission” has been developed, where unacknowledged transmission is selectively performed on frame basis depending on the traffic characteristics of a frame and network error conditions. Its validity that energy efficiency is achievable without sacrificing performance has been evaluated on the real HTTP traffics.

The main goal of this study is to provide an unacknowledged transmission transparently in current 802.11 networks. Application-level and transport-level issues such as friendliness to unacknowledged transmission are not considered as a major topic because they might make unacknowledged transmission too complicated to believe its usefulness. A simple and transparent implementation might become questionable and the conditions good for unacknowledged transmission might become vague and skeptical. It should be validated first that unacknowledged transmission is practically implementable and has potential to use. To address this point, this paper focuses on a simple implementation of selective unacknowledged transmission at the link-level and its performance characteristics. Scenarios of applications and network environments where unacknowledged transmission is appropriate should be studied in detail at the next step.

The paper is organized as follows. First, related works are discussed briefly in the next section. Section 3 describes the concept of unacknowledged unicasting, the key method, and an implementation. Section 4 discusses its energy and performance characteristics in comparison with 802.11 through various measurements and simulations. Selective unacknowledged unicast transmission is proposed from analysis of which traffic conditions are favorable to removing ACKs, and its effectiveness is evaluated with simulation based on a model of real HTTP traffics in Section 5. We then conclude with contributions and future work generated by this research in Section 6.

2. Related Work

For multimedia communication applications in 802.11 networks several researches have been proposed ideas to reduce ACKs in consecutive frames. One of them is “block ACK” of 802.11e [2], which is merged later into the new base standard named IEEE 802.11-2007 [3]. On bulk transmission, a block of data frames are transmitted consecutively from the originator to the recipient. If some bandwidth allocation is performed in advance, the whole block can be handled as one unit and then ACKs can also be aggregated into one ACK. That is, instead of transmitting an ACK for each data frame, one ACK is transmitted only at the end, as a block ACK, to notify the reception status of the whole block. 802.11-2007 also adds “no ACK” policy in QoS facility. When no ACK policy is used in QoS enhanced networks, the recipient takes no action upon receipt of a frame.

Both block ACK and no ACK schemes are newly added in the standard for QoS support and are not mandatory implementation in 802.11. Due to handshaking overhead to set up bandwidth reservation block ACK has meaning only for pre-scheduled bulk data transmission. No ACK scheme is still at the stage of definition and there are little mentions about when or how to use it.

Some 802.11 chip sets such as *Atheros* provide an extension of no-ACK mode where the receiver do not send an ACK frame. With ACK-waiting timer at the sender, it is possible to skip waiting an ACK for retransmission. Combining these features, an unacknowledged unicast transmission is implementable and there exist some reports on this, for example, Mad WiFi driver studies. However, to work with this implementation, both the sender and receiver should agree on ACK mode. It means that switching between ACK mode and no-ACK mode

requires pre-arrangement and applying unacknowledged transmission on frame basis is not possible. D-proxy [4] has covered issues in utilizing this unacknowledged transmission and developed a TCP connection relay to avoid potential performance degradation.

Most studies on energy saving in 802.11 concern sleep to save energy consumed during idle time. Some cover utilizing 802.11 PSM [5], and others introduce a new access control, not compatible with 802.11 [6][7]. In contrast, our study saves energy consumed for active communication, and it is compatible with 802.11.

Reducing ACKs has been also studied in some reliable multicast protocols [8] and multimedia protocols [9], but they are ACK messages at higher layers rather than the link layer like 802.11, which means ACK frames at the link level are still transmitted. The unacknowledged transmission of 802.15.4 mentioned already is close to ours. RLC(Radio Link Control) of UMTS also has unacknowledged mode [10]. However, their target systems are not a wireless LAN, and the usage and effectiveness have been rarely studied.

In summary, our study is novel in that 802.11 unicasting is transparently implemented with 1-to-1 802.11 multicasting. It is a practical approach that can be easily applied to currently working wireless LAN systems. Our experimental results could provide general insights on practicality of unacknowledged transmission.

3. A Simple Unacknowledged Unicast Transmission Utilizing Multicast Transmission

3.1 The Idea of Unacknowledged Unicast Transmission

In 802.11 DCF, the default mode in wireless LAN, every packet destined to a node is transmitted in unicast transmission as shown in Fig. 1. To reduce the control overhead, RTS/CTS exchange may be opted out by adjusting RTS/CTS threshold as in the second diagram. In both cases the receiver sends an ACK frame to the sender to confirm the reception of a frame, and it is the key process for link-level reliability. On the other hand, multicast/broadcast transmission has no ACK after a data frame. Reliability is not provided at the link-level, and hence error recovery should be done, if necessary, by the higher layer protocol.

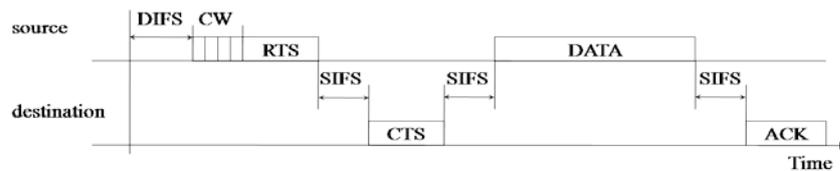
An ACK takes transmission of 14 byte ACK frame plus 192 μ s preamble. The relative cost varies depending on the size of the precedent data frame. For example, in 11 Mbps 802.11b with 28 byte overhead, it could be approximately 14.6% ($\frac{192+(14+28)*8/11}{192+(1500+28)*8/11+192+(14+28)*8/11}$) of 1500 byte data transmission, and 43.8% of 100 byte data transmission. Considering 50% of traffic on LAN is less than 100 byte [11], the overhead of ACK may not be negligible.

This study proposes a novel method to transmit a unicast packet without accompanying an ACK in the current 802.11 system. Since ACK is sent by a receiver's NIC (Network Interface Card), it cannot be removed as long as the unicast MAC address of the receiver is used. If we can do a trick to use a multicast address destined only to the receiver, the unicast packet is encapsulated into a multicast frame and then transmitted without any ACK in the multicast manner. From point of view outside 802.11, unicasting without ACK is achieved; we call it *unacknowledged unicasting*.

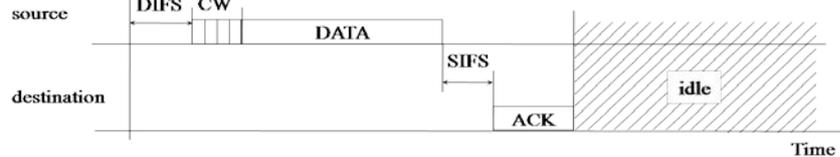
If a unicast packet is transmitted in the multicast/broadcast manner, the energy and bandwidth for the ACK can be saved on condition that the occurrence of errors is handled properly. For some applications that require reliable data transfer, they use a reliable transport

protocol such as TCP or have their own error recovery function. Thus, lack of error control at the link layer has effect not on reliability itself but on resulting performance because recovery may take longer time. For other applications such as multimedia or sensor applications that tolerate data loss to a certain extent, the side-effect is negligible unless link error happens over the limit.

•Unicasting with RTS/CTS exchange



•Unicasting



• Multicasting or Broadcasting



Fig. 1. 802.11 DCF Timing Diagram

It is worthwhile to note that transmitting a frame in multicast does not cause extra energy consumption at neighbor nodes. Even if data is transmitted in unicast frame, neighbor nodes uselessly consume energy to receive and discard the frame. It is known as *overhearing problem*. Therefore, our unacknowledged unicast is not harmful to neighbor nodes.

The effective benefit of unacknowledged unicast is a complicate issue. Although the size of an ACK frame is fixed, its relative weight varies depending on the size of the precedent data frame as mentioned above. The frame error rate, the consequence of not only wireless link condition but also the size of data frame, would be a major factor. The cases when unacknowledged unicast is beneficial may include

- when the application allows some data loss,
- when the link is in good quality and not crowded,
- when the data size is small,
- when the end-to-end delay is short.

Our reasoning partially coincides with rationale that 802.15.4 introduces *unacknowledged transmission mode*.

3.2 802 MAC Address Transformation

The key problem of unacknowledged unicasting is how to decide a multicast address representing the unicast MAC address of a receiver node. In order to apply our scheme to the current 802.11 system transparently, the method should satisfy the following two requirements:

1. Bidirectional transformation

To hide the fact that data is transferred by multicast transmission, the receiver should be able to recognize that the multicast frame is originally unicasted to itself and restore the original unicast frame for next processing. It implies that the unicast address can be driven back from the multicast address, which is transformed from the unicast address at the sender.

2. No disturbance with the existing 802 multicasting

IEEE 802 already supports several multicast services such as IP multicast. If unacknowledged unicasting uses one of the addresses assigned to those services, it may disturb the receiver and other nodes in the LAN.

In conclusion, we need an address transformation scheme to make 1-to-1 mapping between an 802 unicast MAC address and an 802 multicast address.

Every IEEE 802 device has a unique 48-bit address, which is used in unicasting to the device. The first three octets known as the *Organizationally Unique Identifier (OUI)* identify the organization that issued the address. IEEE 802 also defines a multicast address where the first bit in transit, called the *multicast bit*, is set to 1. (The multicast bit is the least significant bit of the most significant byte in binary representation.)

From the current assignments of OUIs [12], only several blocks are assigned in both unicast and multicast address spaces, which means it is possible in space-wise to map each unicast address into a unique multicast address. Fortunately, a simple address transformation rule could be found as follows:

$$\mathit{unique_multicast_address} = \mathit{unicast_address} + 00010001 \overset{\leftarrow 40}{0 \dots 0}_{(2)} \quad (1)$$

The 8th-bit 1 makes an address multicast, and the 4th-bit 1 directs it into one of the empty blocks without conflict with the existing multicast address assignment. Fig. 2 illustrates the address spaces before and after transformation. The first column shows the unicast address blocks currently assigned by IEEE. Since the address transformation changes the OUI portion of an address, each address block is changed as shown in the second column. They are all well shifted into the empty multicast address blocks without any conflict with other address blocks already assigned as in the third column.

3.3 Implementation of Unacknowledged Unicasting

The place where the address transformation is performed is at the point of moving data to/from the NIC, in principle. A few lines of code are enough to implement unacknowledged unicasting. Fig. 3 shows the modified code segments of the Linux networking kernel in pseudo code. The comments describe the code sections around where the modification is made. The codes in italic are added for unacknowledged unicasting whereas the codes in normal are the original code.

Assigned OUI	Transformed OUI	Assigned Multicast MAC addr.
00:00:00 ~00:E6:D3	11:00:00 ~11:E6:D3	01:00:0C ~01:DD:01
02:07:01 ~02:E6:D3	13:07:01 ~13:E6:D3	03:00:00 ~03:00:40
04:0A:E0 ~04:E0:C4	15:0A:E0 ~15:E0:C4	09:00:02 ~09:00:87
08:00:01 ~08:BB:CC	19:00:01 ~19:BB:CC	0D:1E:15
10:00:5A ~10:00:E8	21:00:5A ~21:00:E8	33:33:00 ~33:33:FF
A0:6A:00 ~AA:00:04	B1:6A:00 ~BB:00:04	AB:00:00 ~AB:00:04
		CF:00:00

Fig. 2. Address Space Maps Before and After Address Transformation

Modified Transmit Code (Sender)

```

/* eth_addr has unicast MAC address of the destination, i.e.
   eth_addr = matched address of ARP (IP addr. resolution) */
if unack_flag then
  eth_addr = eth_addr + 11000000000016
endif
dest.eth_addr = eth_addr
/* ... output the frame to the NIC ... */

```

Modified Receive Code (Receiver)

```

/* ... when a frame from the NIC ... */
if (dest.eth_addr has multicast address bit) then
  if ((dest.eth_addr - 11000000000016)
      == device.eth_addr) then
    dest.eth_addr_type = unicast
    dest.eth_addr = device.eth_addr
  else
    dest.eth_addr_type = multicast
  endif
endif
endif

```

Fig. 3. Implementation Code of Unacknowledged Unicasting

The transformation rule and its inverse are coded in straightforward manner, respectively. At the sender the MAC address of the unicast receiver (*eth_addr*) is transformed just before the destination address of a frame is finalized. At the receiver if a received frame is multicast, it is tested whether the destination address is transformed from the receiver's MAC address. If

it is, the address and the type are recovered, and afterwards the frame is treated as a normal unicast frame. Otherwise, usual processing for multicast is followed.

The above implementation may cause incompatibility problem when a sender uses unacknowledged unicasting but the receiver cannot understand unacknowledged unicasting; a frame destined to an unknown multicast address is discarded. It can be solved by setting off the unacknowledged flag (`unack_flag`) for a node. There are many strategies on how to decide `unack_flag`. For example, if no frame has been received from a node with the transformed destination address for a period of time, its `unack_flag` is set to false. In the test implementation, it is simply assumed that a receiver can understand unacknowledged unicasting.

It is worthwhile to note that an unacknowledged node understands 802.11 unicasting. Since the implementation at a receiver only extends the case of multicast address, a unicast frame is processed as it used be. Also the original 802.11 multicast is never affected. In summary, there would be no compatibility problem with 802.11.

4. Performance Implications of Unacknowledged Transmission

Extensive experiments have been conducted to analyze how much unacknowledged unicasting has effect on energy consumption and performance. Some results have been obtained by using ns-2 simulator [13], and others have been measured by experimenting on an implementation where a Linksys WAP11 AP and Linux notebooks are networked in 802.11b.

4.1 Measured Performances in Typical 802.11 Environments

In order to check the feasibility of unacknowledged unicasting, we have measured performances while running some well-known applications in a typical 802.11 environment. First, VoIP was tested as an example of multimedia applications using UDP. While calling between a PSTN phone and a Skype phone on a wireless device, traffics on the wireless link were analyzed using a packet capture tool. **Table 1** shows a summarized result of the Skype data traffics while 100 second-long music was played at the PSTN phone. For each of the normal 802.11 and unacknowledged unicasting, the numbers of voice data frames and ACK frames were counted from the traces captured at the third monitoring station. The energy consumption at the Skype phone was calculated based on a typical power consumption rate (Tx Power = 0.660W, Rx Power = 0.395W) [13].

Table 1. Measured Performance of a Skype Call in Small Room Case

Node Types	Number of Frames		Energy Consumption at Skype Phone (Receiver) (Joule)
	Data	ACK	
802.11	5123	5085	0.279
Unacknowledged Unicasting	5055	0	0.168

As expected, there is no ACK frame in unacknowledged unicasting. Some voice data frames might not be delivered to the Skype phone, but there was no noticeable degradation on voice quality, probably due to the Skype recovery feature. With respect to energy, unacknowledged unicasting consumes about 39.8% less. This big gain comes from the fact that voice data frames are relatively short (mostly 120 bytes), and Tx Power is greater than Rx Power though an ACK frame is 42-byte long only. In summary, for multimedia applications

which are flexible on reliability and hence use UDP, unacknowledged unicasting gives little harm but potential gain.

For reliable communication over TCP, three cases have been set up: an AP and a station are inside a small room, they are at each end of a large office room with many partitions, and a station is outside a room. They are labeled in this paper as *small room*, *large room* and *outside room* case, respectively. The bit error rates, computed from measurement of frame errors, are approximately 1.3×10^{-6} , 3.2×10^{-6} and 5.9×10^{-6} , respectively.

The wireless station downloads a huge file at full speed from a server connected at a wired LAN through the AP. For each case above, the integrity of the downloaded file has been tested and the throughput has been measured. File integrity has been always kept. This implicates that the implementation works correctly and a frame loss could be fully recovered by TCP.

The throughputs are illustrated as in Fig. 4, in comparison with the original 802.11, i.e., *acked unicasting*. In *small room* case the unacknowledged transmission gives higher throughput because the bandwidth for ACK in the original 802.11 can be used for data in the unacknowledged transmission. In *outside room* case the unacknowledged transmission gives lower throughput because first, loss recovery by the end-to-end TCP error control in the unacknowledged transmission takes longer time than by ACK in the 802.11, and second, TCP's congestion control triggered by segment loss in the unacknowledged transmission reduces data sending rates. *large room* case gives in-between results.

This basic experiment shows that unacknowledged unicasting can give comparable throughput to the original 802.11 inside an intranet. As the 802.11 becomes more popular, *outside room* case will be rare and *small room* case will be dominant. Therefore, it would be expected that unacknowledged unicasting will have more potential.

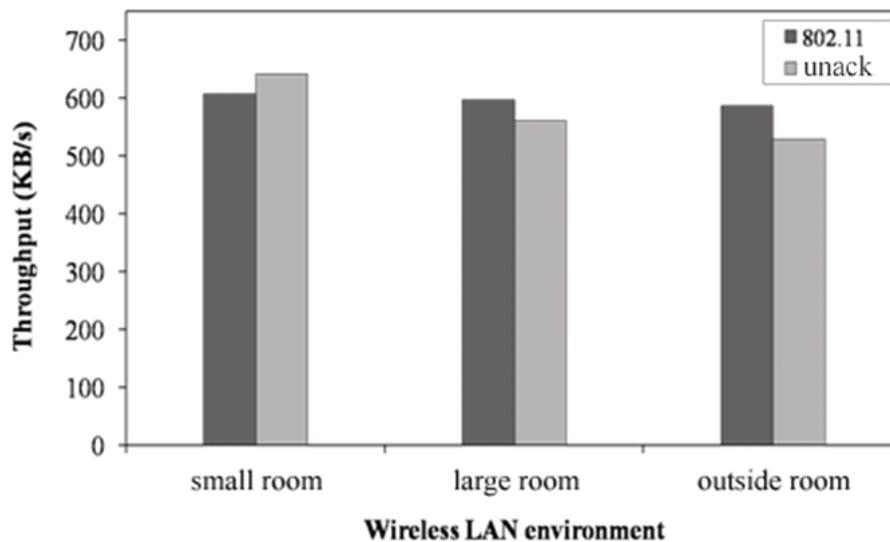


Fig. 4. Measured Throughputs (File Download Speed) for Typical Wireless LAN Environments

4.2 Energy Saving of Unacknowledged Unicasting

To analyze energy characteristics of unacknowledged unicasting, the 802.11 and its energy module [14] of ns-2 have been used, where energy consumption rates are 0.660W for transmitting, 0.395W for receiving, and 0.035W for idling. Link quality and frame size are

important factors on energy consumption. The bit error rate (BER) is usually used to describe a link quality. For a given BER, frame error rate increases as frame size increases, in an approximately linear way [15].

It is simulated that one node downloads a huge file at full speed from the other node via wireless link, by unacknowledged unicasting and 802.11 unicasting, respectively. The BER of the link is varied from 0 to 15×10^{-6} . The frame size of WLAN traffics shows a bimodal distribution where each mean is 100 B (byte) and 1500 B around, respectively. Two cases of frame size have been considered: 100 B for small frames and 1500 B for big frames. Fig. 5 shows the energy efficiency of each method in byte per joule (B/J).

The results indicate that unacknowledged unicasting mostly gives higher energy efficiency than 802.11, up to 41% in 1×10^{-6} BER. Even in 5×10^{-6} BER that is similar to outside room case in the previous measurement, unacknowledged unicasting shows about 38 % better energy efficiency.

However, in higher error condition than 15×10^{-6} BER, which means about 18% frame error rate in 1500 B frame size, the unacknowledged transmission gives worse efficiency. As already mentioned, the performance of unacknowledged unicasting is more sensitive to errors; it will be explained in detail in the next subsection. Unacknowledged unicasting takes much longer times for downloading, most of which are passed in idle state. Energy consumption for idling increases and it causes worse energy efficiency.

In case of 100 B frame, there is very slightly degradation as bit error rate increases, but it is unnoticeable in the figure. It is believed that there are two reasons. First, if the maximum segment size is 100 B, the bottleneck point is not link but TCP. Although higher error rate makes lower the effective bandwidth of link, it has a very little effect on the end-to-end performance. Second, for a small size frame like 100 B, frame error rate is still very low in high bit error rate. The TCP operations are affected infrequently.

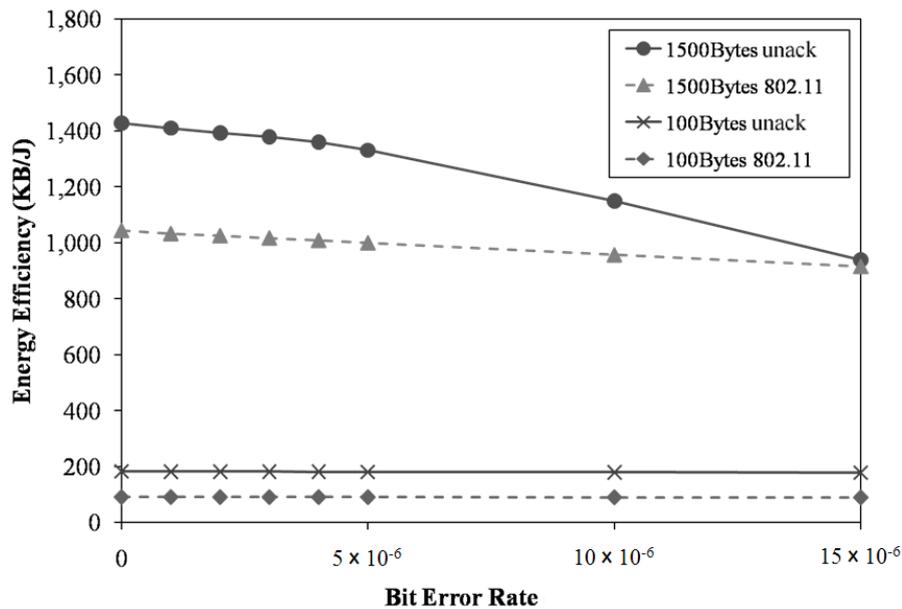


Fig. 5. Energy Efficiency of Unacknowledged Unicasting In One-hop Link

One thing to note is that we have compared energy efficiency for a task, i.e., file downloading. In a real situation, a node may execute a set of communication tasks with user thinking time in the intervals of them. Unless the node switches off the 802.11 interface for thinking time, it keeps consuming idling energy no matter when a communication task is completed. That is, sometimes energy consumption for a given time interval may be a proper metric, and energy consumption during idle time may be factored out in comparison. This issue will be revisited later.

4.3 Effects of Unacknowledged Unicasting on Performance

With respect to energy, unacknowledged unicasting almost always gives better efficiency because it does nothing to improve link quality. In case link quality is poor, this lack of immediate action causes late and miscalculated treatments at the higher layer. As expected conceptually and revealed in the above measurement, its performance might be so worse than 802.11 that the benefit of energy saving is not worthwhile. The question to be addressed here is on which condition the performance of unacknowledged unicasting goes really bad.

First, performances over one-hop 11 Mbps wireless link, which models an intranet, have been evaluated in simulation. A huge file is downloaded at full speed as before. Fig. 6 shows throughputs of unacknowledged unicasting and the original 802.11 unicasting in various bit error rate and frame size. The unacknowledged transmission gives better results except in very high error rates. Bandwidth for ACKs in 802.11 is used for data in unacknowledged unicasting, and it increases throughput when errors are rare. The gain becomes smaller as an error rate goes higher. In comparison to the measurements in Fig. 4, there are slight differences in absolute value but they show consistency. Both results indicate the same characteristic that unacknowledged unicasting responds more sensitively to link error rate.

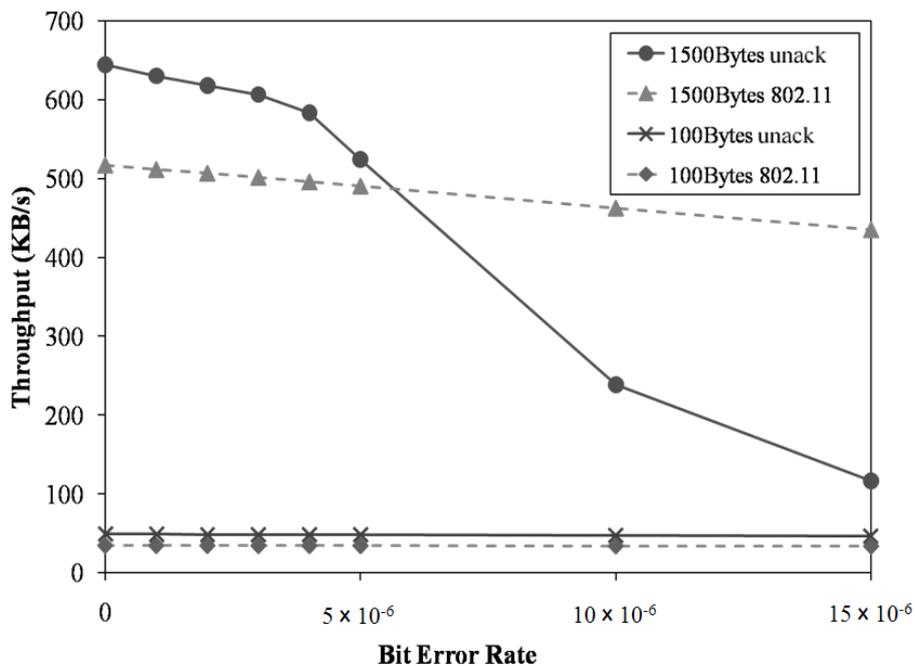


Fig. 6. Throughputs over One-hop Wireless Link

Next, to investigate the performance of Internet access through ADSL, a two-hop path is set up where a relay node and a 10 Mbps wired link with 100 ms delay are appended to the above wireless link. The unacknowledged scheme is applied only to the wireless link, and no error is assumed in the wired link. Fig. 7 illustrates the effect of unacknowledged unicasting on long distance performance. In comparison to the intranet case, the unacknowledged transmission drops sharply as BER increases. It is much beyond the amount of wasted time and bandwidth for error recovery.

TCP considers a packet loss a symptom of congestion and reduces sending rate by decreasing the congestion window. Since unacknowledged unicasting does not recover a frame error, a high link error rate causes lots of packet losses at TCP, and thus the window is decreased too frequently to be up to available capacity. A small congestion window is not a matter in an intranet because the delay-bandwidth product is small. The round-trip time is so small that the window is refilled by a quick feedback before exhausted. However, in Internet with a long round-trip time, it is exhausted before a long delayed feedback arrives, which makes a sender pause to be idle. This type of problems has been studied in so-called wireless TCP [16], and the unacknowledged approach makes it worse. Our simple solution will be presented in Section 5.

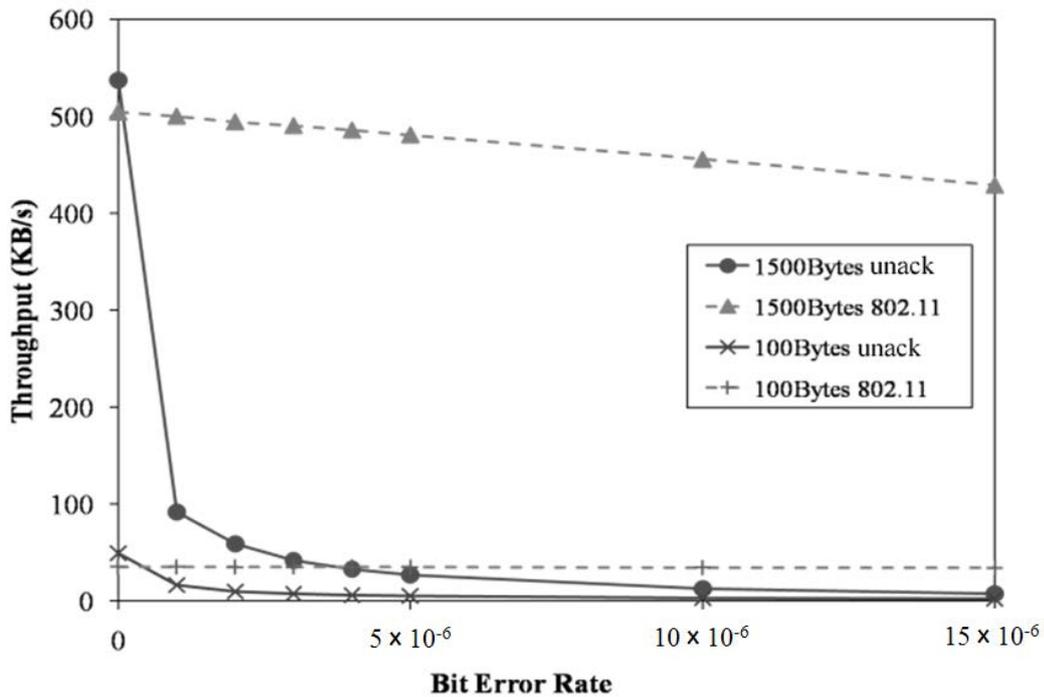


Fig. 7. Throughputs over Two-hop Path with Long Distance Wired Link

4.4 Coexistence of Unacknowledged Unicasting and 802.11

The last point to be addressed in the experiments is how unacknowledged unicasting coexists with 802.11. In other words, the effects of unacknowledged unicasting on 802.11 neighbor nodes have investigated. An AP connects two nodes in almost the same physical condition, where one runs in unacknowledged unicasting and the other does in the original 802.11. While two nodes simultaneously download a huge file at full speed in *small room* condition, their

throughputs have been *measured* as shown in **Table 2**. It also shows the results of when both nodes are in unacknowledged unicasting, and when both are in 802.11.

When both nodes run in the same unicast method, they evenly share the link capacity, where the unacknowledged transmission is slightly higher due to removing ACKs. However, when they run in different methods, the 802.11 node has much higher throughput than the unacknowledged transmission node. It might be because the TCP of the unacknowledged transmission node not immune from a frame error reduces its sending rate, and it gives the 802.11 node a chance to increase the sending rate.

Table 2. Throughputs of Two Nodes Sharing an AP

Node Types	Throughputs
(802.11, 802.11)	(297KB, 294KB)
(802.11, Unack)	(382KB, 245KB)
(Unack, Unack)	(337KB, 342KB)

In order to explore it from microscopic view, their link access patterns are monitored by the third node. Using IOGraph utility of WireShark [17], **Fig. 8** displays the traces of their sending rates for an interval. It also shows when TCP retransmission happens as a tiny box at the bottom. The sending rate of the unacknowledged fluctuates with big drops after TCP retransmissions, but the 802.11 keeps relatively consistent rate with some rises on drops of the unacknowledged transmission. The 802.11 never allows a spacious room like the unacknowledged transmission where the other easily increases its rate. This is the reason why the unacknowledged transmission could not share the half of the link capacity. In case both nodes are unacknowledged transmission mode, they evenly share the link up to full capacity exchanging rises and drops each other.

In summary, unacknowledged unicasting is never harmful but unselfish. It increases link capacity by removing ACKs, and its link access is less competitive. Its presence has nothing but positive effects on 802.11 neighbor nodes.

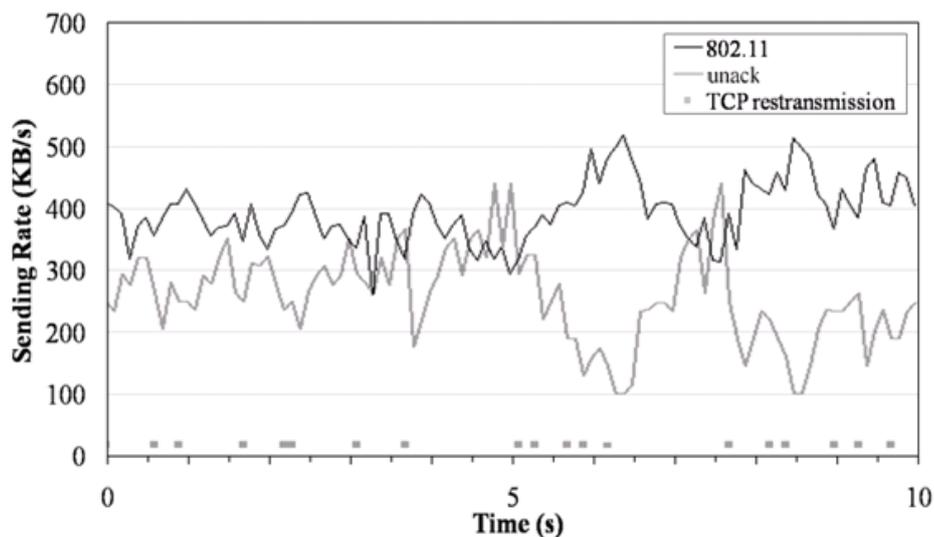


Fig. 8. Trace Graph of Sending Rates at Two Nodes Sharing An AP

5. Selective Unacknowledged Unicasting

5.1 Idea of Selective Unacknowledged Unicasting

The experiments show that unacknowledged unicasting mostly provides energy saving with comparable performance, but in some cases throughputs are downgraded severely that it cannot be compensated by energy saving. The idea of selective unacknowledged unicasting is to dynamically switch unicasting back to the original 802.11 in case bad performance is expected. It is believed that the decision cannot be deterministically correct. Thus, a heuristic selection rule will be made through analytical experiments.

In order to realize selective unacknowledged unicasting, those conditions when unacknowledged unicasting may result in worse performance than the 802.11 should be figured out first. From the experiments it has been found that the following three factors affect the performance of unacknowledged unicasting:

- link error rate
- frame size
- round-trip time (i.e., distance to a receiver)

To classify traffic conditions two representative cases for each factor are considered, and in combination of them eventually eight cases are generated as the representative traffic conditions, as in **Table 3**. For the link error rate, 1×10^{-6} BER is set for the low and 5×10^{-6} BER is for the high, which approximately correspond to *small room* and *outside room* wireless LAN in the previous experiment, respectively. Since it has been well known that the frame size is distributed bimodally in LAN, 100 B for a small frame and 1500 B for a large frame are set. For the round-trip time, 0.01 ms is set for the short distance labeled as *Intranet*, and 100 ms is for the long distance labeled as *Internet*.

Table 3. Analysis of Cases When Unacknowledged Unicasting is Adequate

	Error Rate	Frame Size	Location of Receiver	Throughput	Energy Efficiency
				(Unack-802.11)/802.11	
1)	Low	Small	Intranet	40.64%	106.23%
2)	Low	Small	Internet	-51.97%	63.64%
3)	Low	Large	Intranet	23.25%	41.93%
4)	Low	Large	Internet	-81.50%	-7.80%
5)	High	Small	Intranet	39.51%	105.89%
6)	High	Small	Internet	-84.29%	-0.63%
7)	High	Large	Intranet	10.06%	38.53%
8)	High	Large	Internet	-94.35%	-52.63%

For each case, the same full-speed downloading has been simulated, and throughput and energy efficiency are compared. The metric is relative difference with 802.11 as a basis. It is interpreted that for case 2) unacknowledged unicasting gives 51.97% less throughput but 63.64% more energy efficiency. The decision which cases are adequate for unacknowledged unicasting depends on relative importance between throughput and energy. Under a simple assumption that energy and throughput have the same weight, unacknowledged unicasting is worthwhile to select if throughput plus energy efficiency is positive. The corresponding cases are 1), 2), 3), 5), and 7). These cases can be represented as the following logical expression of

the traffic conditions where `unack_flag` is a boolean variable in the implementation in Fig. 3.

$$unack_flag = Intranet + Low_error \cdot Small_size \quad (2)$$

The next issue is how to implement it in a real system. When an outstanding frame is ready to send to the NIC, the above three factors should be known to determine which case it falls in. Knowing the frame size is trivial. The link error rate is also easily obtainable because almost all LAN drivers of modern operating systems maintain an average frame error rate for diagnosis utilities. For example, Linux system provides `net_device_stats` structure.

It seems non-trivial at the frame-level to know the round-trip time or the final destination. However, whether the traffic goes to an intranet or Internet is still predictable. If the receiver is inside an intranet, the frame is usually delivered directly. However, if the receiver is outside an intranet, the frame should be relayed by a router, and thus the destination address should be the router's MAC address. Therefore, it is somewhat conservative but safe that a frame is considered as Internet traffic if its destination MAC address designates a router.

Implementation of selective unacknowledged unicasting can be done easily by inserting a few lines of code to unacknowledged unicasting given in Fig. 3. First, three traffic conditions are determined as mentioned above, and then unacknowledged flag is calculated as (2) before the address transformation is performed.

5.2 Performance of Selective Unacknowledged Unicast Transmission in Real HTTP Traffics

The practical effect of selective unacknowledged unicast transmission depends on how real traffics are distributed over the above cases. What is most concerned is the real distribution of frame size and location of destination. We have built a traffic generator on *ns* based on a well-known model of real HTTP traffics [18]. The traffic generator emits a list of the parameters of a traffic including frame size and location (intranet/internet). Depending on the location, the client sends a request to a local server on a local link or to a remote server over a long distance link. The corresponding server replies with a message whose size is also determined based on the model. The client pauses to think for amount of time generated, and then the next request is generated. The simulation environment parameters are the same as the other simulations; the RTT of the local and remote server are 0.01msec and 100msec, respectively, and the bandwidth of the wired link is 10 Mbps. Simulation has been made for fifty thousand traffics, which approximately cover network activity for one week.

First, response times for HTTP requests are analyzed for performance comparison because throughput is distorted by intermittent user thinking time. Fig. 9 shows the average response time of each unicast method in various BER. Basically, selective unacknowledged unicasting goes between unacknowledged unicasting and 802.11 because it mixes them case by case. The point to notice is the results at high BER. It is much closer to the 802.11 than the unacknowledged transmission, only 17% higher than the 802.11 at 15×10^{-6} BER. It is conformed to our intention to avoid a severe performance drop of unacknowledged unicasting in a poor quality link. The other point is when BER is 5×10^{-6} , which is the threshold of the *high* error rate. From that point, a small message to Internet in selective unacknowledged unicasting is transmitted in 802.11 manner, and thus its result goes closer to the 802.11.

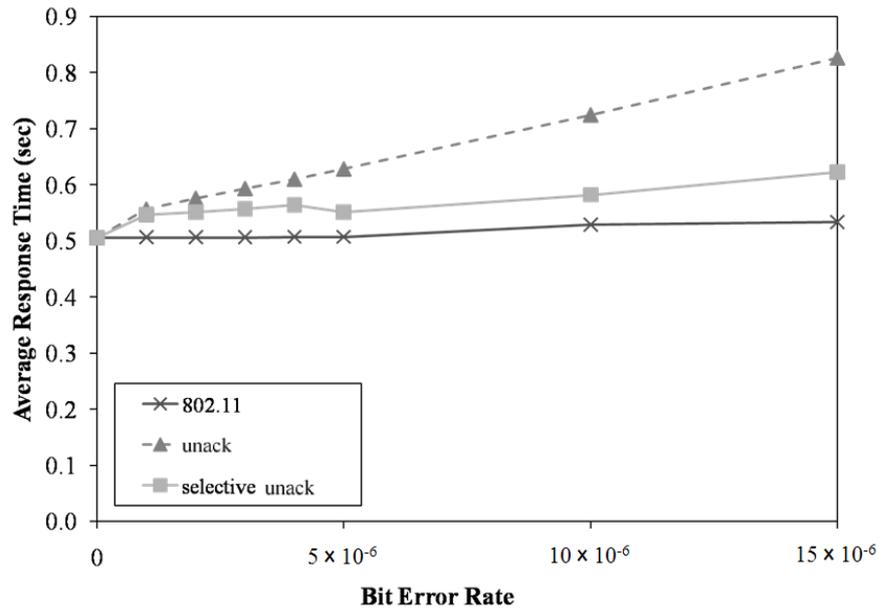


Fig. 9. Average Response Time for HTTP Requests

With respect to energy saving, energy consumptions in joule (J) are compared instead of energy efficiency due to the same reason of thinking time. A total energy consumed to process a given amount of communication tasks has more meaning. The result could be interpreted as energy consumption for one-week network activity. **Fig. 10** shows total energy consumptions of each unicast method. Again, the selective unacknowledged unicasting is placed between the others; it consumes about 15% to 25% less energy than the 802.11. Based on these result, it could be concluded that selective unacknowledged unicasting with a simple rule achieves the goal of energy saving with comparable performance in real usage environment.

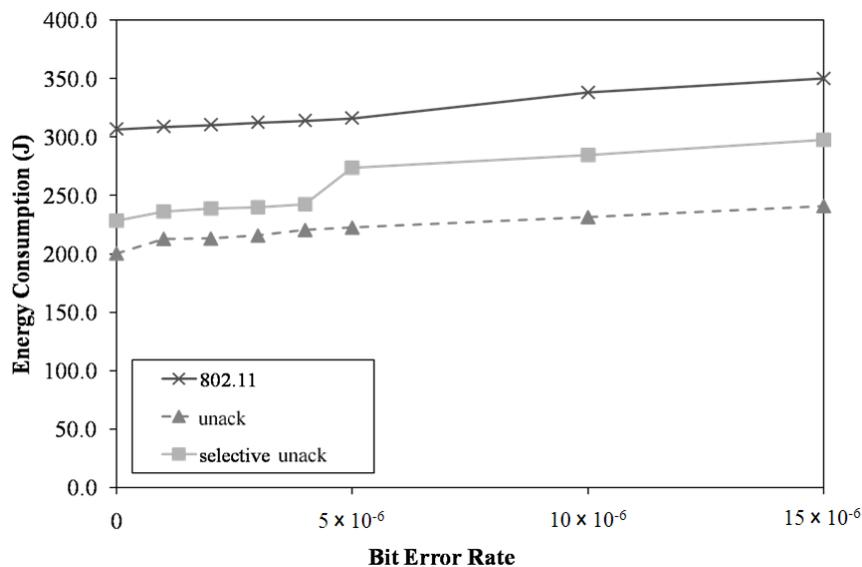


Fig. 10. Total Energy Consumption for HTTP Traffics

6. Conclusions and Future work

As 802.11 becomes very popular, most private wireless LANs serve only a few nodes in a small area. In this type of wireless LAN environment, frame errors are rare and thus an ACK after a unicast transmission may have little meaning, although it consumes energy and bandwidth. This paper introduced a novel and simple method to remove an ACK in unicast transmission. The key idea is to transparently transform a unicast frame to a multicast frame which does not accompany an ACK. Practicality is also a strong point because it is easily implementable and compatible with working 802.11 systems.

This unacknowledged approach makes unicasting more efficient but less strict. Extensive experiments have shown that energy saving is achievable with comparable performance except some traffic cases. Selective unacknowledged unicasting has been proposed to apply removing ACKs only when it is expected to be profitable.

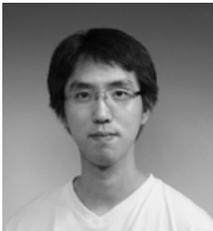
The experiment results give some insight on when and how to utilize unacknowledged transmission. In Skype, a combination of an error-tolerant application and UDP, unacknowledged transmission is so useful to consume only one-half energy without sacrificing quality. In a file transfer application, a pair of a reliable application and TCP, it degrades performance when the corresponding node is remote through a poor link. More experiments with various types of applications and network environments would show clearer guides on which scenarios unacknowledged transmission is useful. Then rules for selective unacknowledged transmission will be extended, including application and transport protocol issues.

An area that needs further study is a correlation with TCP congestion control where a packet loss is considered as a signal to decrease sending rate. Many works have been done to adapt TCP to error-prone wireless networks, but a more aggressive method will be needed in order that unacknowledged unicasting, which let errors happen, is effective in long distance communications.

References

- [1] IEEE Standard for Information Technology Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs), IEEE Standard 802.15.4, 2006.
- [2] IEEE Standard for Information Technology Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements, IEEE Standard 802.11e, 2005.
- [3] IEEE Standard for Information Technology Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, IEEE Standard 802.11, 2007
- [4] David Murray, Terry Koziniec and Michael Dixon, "Solving Ack inefficiencies in 802.11 networks," in *Proc. of IEEE Conf. on Internet Multimedia Services Architecture and Applications*, pp.1-6, 2009. [Article \(CrossRef Link\)](#)
- [5] Ronny Krashinsky and Hari Balakrishnan, "Minimizing Energy for Wireless Web Access with Bounded Slowdown," in *Proc. of ACM Conf. on Mobile Computing and Networking*, pp.119-130, 2002. [Article \(CrossRef Link\)](#)
- [6] Wei. Ye, John. Heidemann and Deborah Estrin, "An Energy-efficient MAC protocol for wireless sensor networks," in *Proc. of IEEE Conf. Computer Communications*, pp.1567-1576, 2002. [Article \(CrossRef Link\)](#)

- [7] Amre El Hoiydi, Jean Dominique Decotignie, Christian C Enz and E Le Rouxz, "WiseMAC: An Ultra Low Power MAC Protocol for the WiseNET Wireless Sensor Network," in *Proc. of ACM Conf. on Embedded Networked Sensor Systems*, pp.302-303, 2003. [Article \(CrossRef Link\)](#)
- [8] Sally Floyd, Van Jacobson, Ching Gung Liu, Steven McCanne and Lixia Zhang, "A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing," *IEEE/ACM Networking*, vol.5, no.6, pp.784-803, 1997. [Article \(CrossRef Link\)](#)
- [9] "Wi-Fi Multimedia™," [Online]. Available: http://www.wi-fi.org/knowledge_center/wmm/
- [10] 3GPP TS 25.322 V5.13.0, 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Radio Link Control (RLC) protocol specification (Release 5).
- [11] "Academic Research Group," [Online]. Available: <https://research.sprintlabs.com>
- [12] "IEEE OUI and Company_id Assignments," [Online]. Available: <http://standards.ieee.org/regauth/oui/index.shtml>
- [13] Kevin Fall and Kannan Varadhan, "The ns Manual," The VINT Project.
- [14] "Energy Model in ns," [Online]. Available: <http://www.isi.edu/nsnam/ns/doc/node224.html>
- [15] Larry L. Peterson and Bruce S. Davie, "Computer Networks a Systems Approach, 3rd ed.," *Morgan Kaufmann*, pp. 92-93, 2003.
- [16] Mario Gerla, M. Y. Sanadidi, Ren Wang and Anrea Zanella, "TCP Westwood: Congestion Window Control Using Bandwidth Estimation," in *Proc. of IEEE Conf. on Global Telecommunications*, pp.1698-1702, 2001. [Article \(CrossRef Link\)](#)
- [17] "The IO Graphs window," [Online]. Available: http://www.wireshark.org/docs/wsug_html_chunked/ChStatIOGraphs.html
- [18] C. R. Simpson, Jr., D. Reddy and G. F. Riley, "Empirical Models of TCP and UDP End-User Network Traffic from NETI@home Data Analysis," in *Proc. of 20th Principles of Advanced and Distributed Simulation Workshop*, pp.166-174, 2006. [Article \(CrossRef Link\)](#)



Hyun Yang received the B.S. and M.S. degrees in computer engineering from Chung-Ang University, Seoul, Korea in 2006 and 2008, respectively. He is currently a Ph.D. Student in the School of Computer Science and Engineering, Chung-Ang University, Seoul, Korea. His research interests include wireless sensor networks, reliable wireless multicast, and energy-efficient communication.



Jin Seok Yun received the B.S. and M.S. degrees in computer engineering from Chung-Ang University, Seoul, Korea in 2007 and 2009. His research interests include cognitive radio, wireless interference, and energy-efficient communication.



Jun Seok Oh received the B.S. and M.S. degrees in computer engineering from Chung-Ang University, Seoul, Korea in 2008 and 2010. His research interests include cognitive radio and energy-efficient communication.



Chang Y. Park (M'93) received the B.S. and M.S. degrees in computer engineering from Seoul National University, Seoul, Korea in 1984 and 1986, respectively, and received the Ph.D. in computer science from the University of Washington, Seattle in 1992. He is currently a Professor in the School of Computer Science and Engineering, Chung-Ang University, Seoul, Korea. His research interests include computer networks, and real-time systems.