Queen's University
Belfast

Implementation of Selective Packet Destruction on Wireless Open-Access Research Platform

Published in:
2013 Asilomar Conference on Signals, Systems and Computers

Document Version:
Early version, also known as pre-print

Queen's University Belfast - Research Portal:
Link to publication record in Queen's University Belfast Research Portal

Implementation of selective packet destruction on wireless open-access research platform

S Hughes[1], B Zhou[1], R Woods[1] and A Marshall[2]
[1]ECIT Institute, Queen's University Belfast,
Queen's Road, Belfast, BT3 9DT, UK
[2]Department of Electrical Engineering and Electronics,
University of Liverpool, L69 3GJ, UK

Abstract—**Interesting wireless networking scenarios exist wherein network services must be guaranteed in a dynamic fashion for some priority users. For example, in disaster recovery, members need to be able to quickly block other users in order to gain sole use of the radio channel. As it is not always feasible to physically switch off other users, we propose a new approach, termed selective packet destruction (SPD) to ensure service for priority users. A testbed for SPD has been created, based on the Rice University Wireless open-Access Research Platform and been used to examine the feasibility of our approach. Results from the testbed are presented to demonstrate the feasibility of SPD and show how a balance between performance and acknowledgement destruction rate can be achieved. A 90% reduction in TCP & UDP traffic is achieved for a 75% MAC ACK destruction rate.**

Index terms: Network Security, Implementation, Denial of Service, Network Management

I. INTRODUCTION

Guaranteed communication services are important for some critical applications such as disaster recovery, law enforcement, and battle field communications etc. However, it is especially difficult to guarantee this, as radio is often the only means of communication in these scenarios. Typically, multiple wireless networks may exist in the same geographical area competing for bandwidth on the same channel. A network administrator may be able to limit the users on the wireless network for which they are responsible, but not on other wireless networks beyond its control. If the administrator tries to block a channel, then generally no-one can communicate over it, including the priority users. We propose a selective packet destruction (SPD) solution to address this dilemma, which can also be used as either an attack tool [1-2], or a network management tool for network regulation.

The rest of the paper is organised as follows. Section 2 presents a brief review of the existing work in this area. Section 3 elaborates the details of selective packet destruction attack and its applications; this is followed by a description of the SPD implementation in Section 4 and the experiments in Section 5. The conclusions are given in Section 6.

II. RELATED WORK

Selective packet destruction can be considered as a type of jamming and has been used as an effective Denial of Service (DoS) attack in various scenarios. We can categorise jamming attacks into two basic strategies: proactive and reactive. Proactive jamming constantly blocks one or more channels by transmitting noises or packets continuously, i.e. permanently occupying the target channels. The disadvantages of this type of attack are that it is energy demanding and it can be easily detected. Reactive jamming on the other hand, does not constantly transmit. Instead, it transmits only when needed, based on its own discretion and attack strategies and therefore is considered to be more intelligent. With advanced strategies, it can be designed to be energy efficient and hard to detect. Patel et al. [1] carried out a survey on attacks against wireless sensor networks including jamming attacks and countermeasures. Pelechrinis et al. [8] reviewed jamming attacks in wireless networks.

A reactive jamming attack, Jamming ACK (JACK), was investigated by Zhang et al. [1]. JACK targets IEEE 802.11 DCF networks wherein an attacker transmits wireless signals to jam (corrupt) ACK messages in order to force the data transmitter to retransmit the data packet. They demonstrated, by network simulations, that a wireless network could be easily disrupted by the JACK attack.

In order to mitigate the adverse effects of jamming attacks, various countermeasures have been proposed. A counter-measure, termed Extended NAV (ENAV), was proposed to inhibit the above mentioned JACK attack. It extends the length of an ACK timeout to let the ACK be transmitted at a random time slot to reduce the chance of it being destroyed by the JACK attack. While this scheme improves the resilience, it reduces the efficiency of the MAC protocol and may have implications for other aspects. For example, the ENAV mechanism may be exploited by a malicious attacker or a selfish node to intentionally change the ENAV in order to disrupt the network or to gain unfair access to the network resource.

In [1], a countermeasure, Random Backoff Control (RBC), against malicious behaviour in WLANs was proposed. The aim of RBC is to assure network resources are fairly distributed among network users by regulating the backoff selection process.

In [1], a Trust Management Framework (TMF) using grey theory and Fuzzy sets was proposed to assess the trustiness of node behaviour. Since multiple network parameters are used for assessment, the TMF is claimed to have high efficiency in detection of misbehaviours and can be used in conjunction with other regulation algorithms.

The contribution of this paper is in the practical implementation of a reactive jamming attack called the SPD which acts against the MAC ACKs, thus validating the practical feasibility of such an attack. We also propose a novel application of this type of attack to regulate network usage. To our best knowledge, this is the first time that such an SDP attack has actually been implemented and demonstrated on a real application.

III. SELECTIVE PACKET DESTRUCTION (SPD)

A. Problem Statement

Consider the scenario illustrated in Fig. 1. An AP provides a network service to users A and B. User A is a priority user while B is an ordinary user. A moderator monitors and regulates the network usage of each user. In an emergency situation, user A should be guaranteed to use the network service while user B can have only limited network usage. This network policy is enforced by the network moderator.

To reduce the bandwidth usage of user B, the Moderator will look for data packets from and to user B. Following each data packet, an MAC ACK frame is expected to be transmitted after a SIFS (short inter-frame space). Based on its observation, the Moderator will decide if it needs to kill that ACK frame. If yes, it will transmit a short packet to destroy the ACK on the fly. Consequently user B will either back off or retransmit the original data packet, as described in section 3.


Fig. 1. Application Scenario

B. Network Model

The network type is one of the following types namely wireless LANs, wireless ad hoc networks, wireless mesh networks, and mobile multi-hop wireless networks. The principle can also be applied to other wireless networks wherein users communicate to each other via broadcast media, adopting the MAC ACK mechanism, i.e. each MAC data frame must be acknowledged by a MAC ACK for the confirmation of successful receipt.

C. SPD Model

There are numerous ways to destroy selected packets on the fly; the most intuitive approach is to transmit noise to collide with packet signals, but the attacker must know the packet ID or MAC address of the sender or receiver of the packet when it is transmitted. The first aspect can be determined through the Carrier Sense (CS) which is available for most wireless chipsets, however, knowledge of the packet ID or MAC address is difficult to acquire in advance, as it cannot be predicted when an unwanted packet will be transmitted. This knowledge has to be extracted while the packet is in transmission. As soon as the packet ID has been extracted, the moderator will determine if it is an unwanted packet and will then transmit some noise signals to destroy the remaining on-going packet transmission.

This presents challenges as the moderator has to respond quickly enough to transmit the noise before the on-going packet transmission finishes. Currently, Wi-Fi data rates have reached 1.3Gbps (IEEE802.11ac) over a single channel using off-the-shelf commercial products which requires less than 1μs to transmit 1000 bits. Thus, the moderator needs to be highly efficient at bit stream decoding, processing, and state transition which represents the process from receiving state to transmitting state when a single radio is used.

An alternative approach involves prediction of the forthcoming packet transmission based on previous packet transmissions. Since the wireless medium is unreliable, most wireless MAC protocols implement a MAC acknowledgement mechanism to confirm the correct reception of a data frame, and/or a virtual carrier sense mechanism to reserve the medium. A MAC data frame will be re-transmitted for a pre-specified maximum number of times before it is dropped if its MAC acknowledgement is never received. One can predict a MAC ACK based on a received MAC data frame and kill it. This may not be as effective as the direct data packet destruction since the actual data frame is not destructed; however it might be energy efficient as it only needs to transmit a very short noise packet whose length is related to the ACK's length.

In the case of virtual carrier sense being used, a data packet is transmitted following the successful exchange of RTS (Request to Send) and CTS (Clear to Send). The transmission sequence is RTS, CTS, DATA, and ACK. One can predict a CTS signal based on a received RTS, a DATA based on a received CTS, and an ACK based on a received DATA. In this scenario, one can choose to destroy either CTS, or DATA, or ACK frames.

In this paper, we chose to use this model to selectively destroy MAC ACKs. We implemented the SPD on the Wireless Open-Access Research Platform (WARP) platform [3] and Fig. 2 illustrates the work flow of the SPD process of our implementation. The moderator keeps monitoring the wireless

media; whenever it receives a packet header it will check the header information against the criteria to see if it matches the pre-defined destruction rules such as frame source /destination addresses, frame type, bandwidth occupied by the source /destination, signal strength and variation, etc. If the criteria are matched, the moderator will then transmit a short noise packet when the medium becomes busy again in order to destroy the expected ACK response from the receiver.
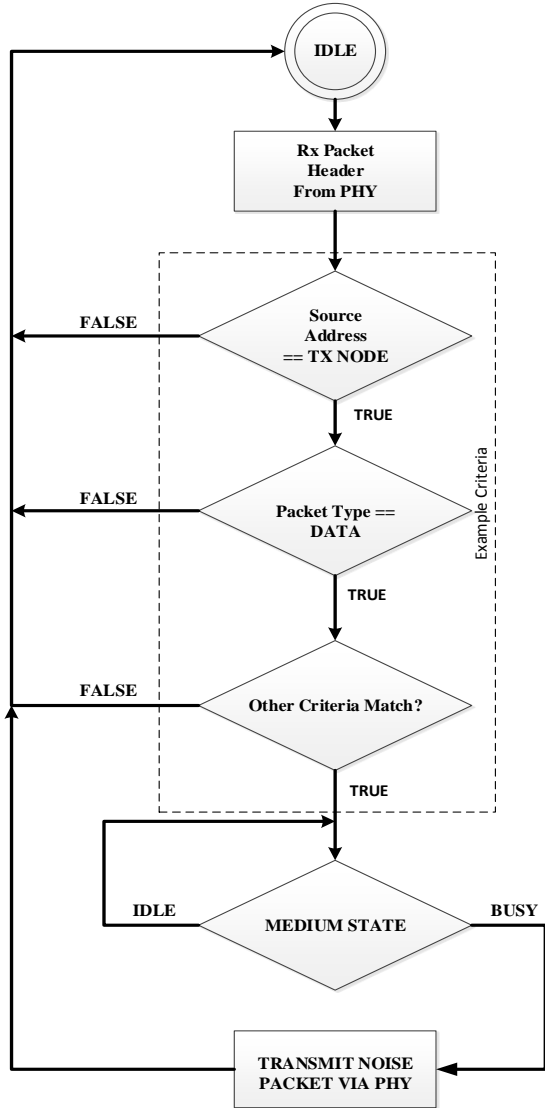

Fig. 2. SPD Work Flow

IV. PROTOTYPING AND EXPERIMENT SETUP

In order to investigate these approaches, the use of the WARP v3 [8] was adopted as it allows the alteration of the physical layer although for these experiments, only the moderator MAC layer was altered. WARP v3 integrates a high performance FPGA including a MicroBlaze processor, two flexible RF interfaces and multiple peripherals to allow rapid customisation of wireless designs. The initial implementation of SPD based on WARP OFDM reference design, is given below.

A. WARP OFDM Reference Design

Fig. 3 shows the WARP OFDM Reference Design which implements a real-time network stack on a WARP node. The design includes a MIMO OFDM physical layer and flexible MAC interface for building custom protocols. The OFDM reference design employed a simplified Carrier-Sense Medium Access (CSMA) protocol which encompasses much of the behaviour of the commercial IEEE 802.11 MAC/PHY chipsets. This algorithm lends itself nicely to a state diagram which can be run on the MicroBlaze of the WARP FPGA; the source code for this is on the WARP website [3].
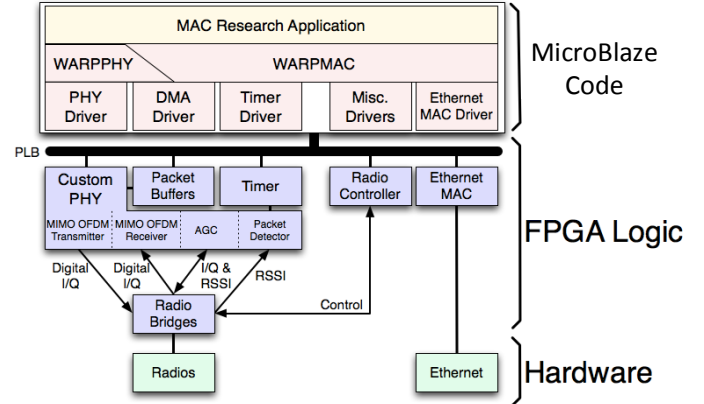

Fig. 3. WARP OFDM Reference Design Structure

For this MAC algorithm, only two types of packets are defined and can be loaded into the MAC frame and understood by the receiver, DATA or ACK. The data packet flag tells the receiving node that a payload is present in this packet and must be delivered via the Ethernet connection to some connected device; it will also send an ACK frame to the transmitting node. If it is an ACK, the node must have been in a timeout waiting on an acknowledgment. It will clear the timeout and return to the idle state.

The WARP MAC employs a retransmission and backoff mechanism which is equivalent to that of the IEEE 802.11 MAC. A data frame will be re-transmitted until it reaches the maximum number of retransmissions, namely n. The default value of n is 8 which can be changed and the maximum back-off window size is $2^{n+4}-1$.

B. Experiment Setup and Prototyping Validation

Implementation involved coding in C to add the functionality of SPD shown in Fig. 2 into the WARP MAC source. As depicted in Fig. 4, a testbed was set up with two PCs and three WARP nodes (A, B and Moderator); each PC is connected via a wired Ethernet connection to a WARP node. Each packet received at the WARP node's Ethernet interface is encapsulated in the WARP's simplified CSMA/CA protocol and transmitted over the air to the other WARP node. Every data frame that is received wirelessly is sent to the WARP node's Ethernet interface, this allows the PCs to communicate via Ethernet over the wireless bridge. Tools used included Iperf
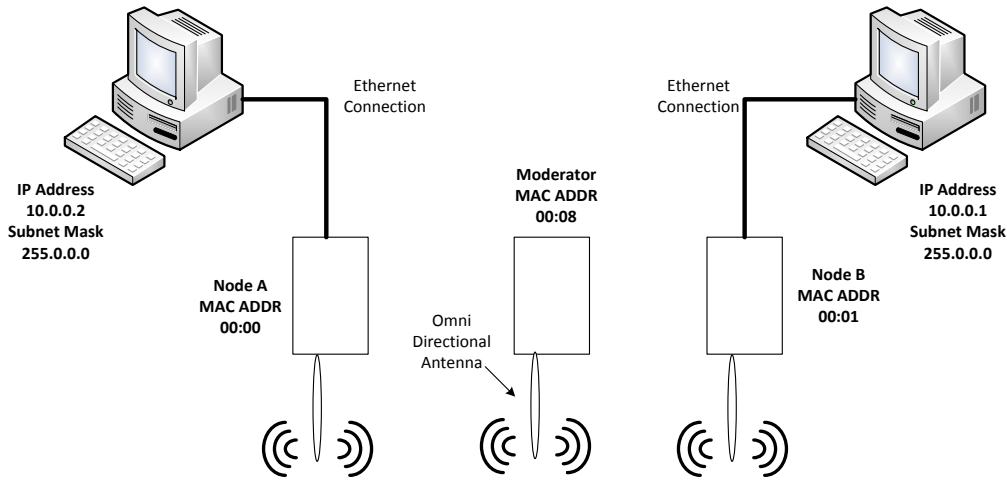
Fig. 4. Experimental Setup

(to generate TCP and UDP traffic), PuTTY (to collect packet delivery statistics), and Wireshark (to record network traffic for further inspection).

The verification of the SPD attack involves two stages. Firstly, the packet statistics are checked to ensure that only the MAC-ACK frames are destroyed i.e. no MAC-DATA frames lost due to attack, and also that the amount of MAC-ACK frames lost due to the attack is equal to the desired attack rate. This is accomplished by processing the statistics of DATA frames sent/received and MAC frames sent/received on Nodes A and B.

The second part of verification is to assess the impact that the attack has on the throughput of the link, with the aim of reducing throughput as the attack rate is increased. The network performance measurement tool *iperf* is used to generate TCP and UDP data streams and report the bandwidth between the PCs. All validating experiments were taken with no moderator active to get a baseline for the performance.

V. EXPERIMENT RESULTS

The testbed uses five pieces of equipment: three WARP v3 boards and two desktop PCs as in Fig. 4. Each PC connects to a WARP board through an Ethernet cable. One WARP board acts as the Moderator. The Moderator constantly monitors the channel and destroys forthcoming ACKs when the criteria are met. Extensive experiments were carried out with various parameters such as different packet sizes and MAC ACK destruction rates. Due to space limitation, we only present part of the experiment results here.

A. Attack Performance

Prototyping faces more challenges than simulation. In a real environment, there are many uncertainties, each of which may affect experiment results. Fig. 5 illustrates the change in MAC ACK loss rate for both TCP and UDP traffic as the attack rate varies. It can be observed that the ACK loss rate shows a strong correlation with attack rate until it reaches 75%. Beyond this

point, the ACK loss rate decreases while the attack rate increases. This might be due to the restraint of the hardware capability, e.g. processing power, timing, etc. In order to destroy a MAC ACK frame, the Moderator needs to detect the packet transmission, match the criteria, and quickly change its state from receiving to transmitting. The SPD matching code is running in a MicroBlaze CPU clocked at 160 MHz which is relatively slow compared with the CPU frequency of modern radio chipsets. The attack performance might be improved via code optimization; however, Fig. 5 has shown that it is feasible to selectively destroy frames based on the WARP v3 platform.


Fig. 5 Attack Rate vs ACK Loss Rate in Two Series of Attacks

B. TCP Performance

Fig. 6 shows the TCP throughput under SPD attacks, which target only MAC ACKs in one direction (Node B - Node A). We can see that the TCP throughput decreases with a strong correlation to the MAC ACK loss rate increasing, which means more bandwidth is released. When the MAC ACK loss rate reaches 75%, the reduced throughput (i.e. 1-TCP throughput) reaches nearly 90%. In this way, a large amount of released bandwidth can be used by priority users. It can also be observed that the TCP packet size change does not have significant impact on the attack performance.

It has been shown that a TCP packet loss rate of 20% can effectively bring down a TCP connection [1]. It is different in the SPD model. SPD kills MAC ACKs with a certain probability but data frames are not directly targeted. So, the TCP packet loss rate is not affected in SPD since TCP DATA and TCP ACK can be sent to the other end without problems. This has been observed in the experiment results. SPD significantly increases the interval between two consecutive TCP segments, which slows down the TCP link.
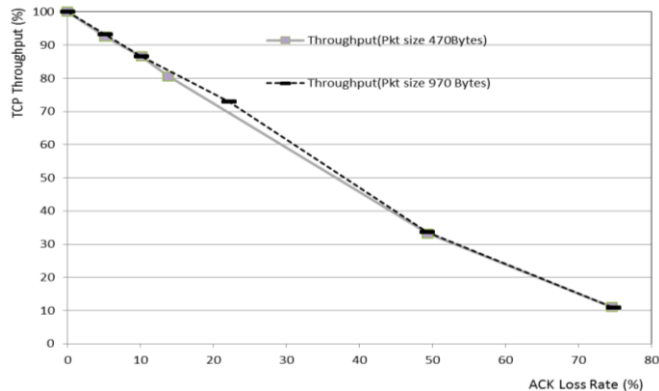

Fig. 6 TCP Throughput versus MAC ACK Loss Rate

C. UDP Performance

Fig. 7 presents the UDP maximum changes under single direction (Node B - Node A) SPD MAC ACK attacks. The performance changes are very similar to that of TCP as shown in Fig. 6.
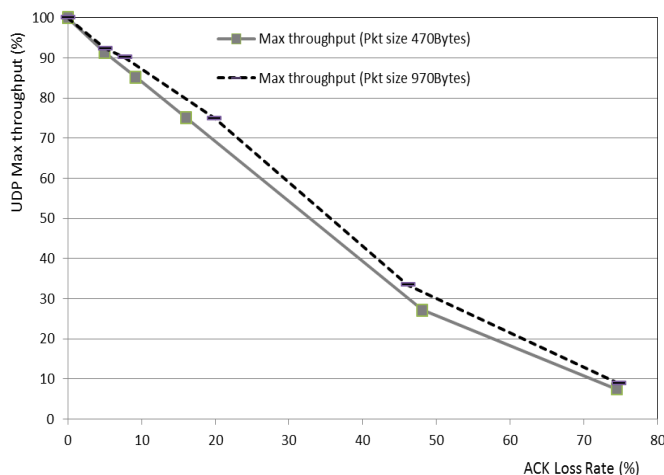

Fig. 7 Maximum UDP Throughput under Attacks

For UDP, delay and jitter are important performance parameters. Since the two communicating nodes are not synchronized, we are not able to measure the delay performance. Fig. 8 shows the jitter changes when MAC ACK loss rate increases. It can be observed that when the ACK loss rate increases to 50%, the jitter rapidly increases. Beyond this point, the number of retransmissions of a UDP packet may vary greatly and this will lead to large uncertainties in the delay between two consecutive UDP packets, i.e. large jitter. The jitter is slightly larger for bigger size packets. This is due to the packet transmission delay contribution to the jitter.

Other experiments show similar effects of SPD attacks and again demonstrate how this type of attacks can effectively reduce a link's bandwidth.
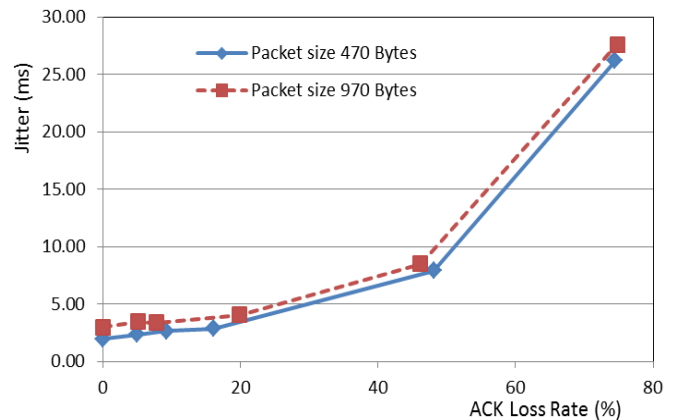

Fig. 8 UDP Jitter under Attacks

VI. CONCLUSIONS AND FUTURE WORK

We have presented and implemented the Selective Packet Destruction (SPD) attack on WARP v3 platform and applied it to regulate network traffic so that priority users can be guaranteed to use network resources in certain circumstances such as in disaster recovery and law enforcement. Experiments have demonstrated the feasibility of implementation of SPD attack, and its effectiveness in regulating network traffic.

Future research directions include a theoretical analysis of SPD efficiency, establishment of the relationship between the attack rate and the actual released bandwidth, and comparison of the efficiencies of different attack strategies.

REFERENCES

[1] B. Zhou et al., "A Random Packet Destruction DoS attack for Wireless Networks", IEEE Int. Conf. on Communications, Beijing, China, May 19-23, 2008, pp. 1658-1662.
[2] M. M. Noor, W. H. Hassan, "Wireless Networks: developments, threats and countermeasures", Int. Journal of Digital Information and Wireless Communications, Vol. 3, pp. 119-134, No. 1, 2013.
[3] M. M. Patel and A Aggarwal, "Security attacks in wireless sensor networks: A survey", Int. Conf. on Intelligent Systems and Signal Processing, 1-2 Mar. 2013, pp.329-333.
[4] K. Pelechrinis et al., "Denial of Service Attacks in Wireless Networks: The Case of Jammers", IEEE Communications Surveys & Tutorials, vol.13, No.2, pp.245-257, May 2011.
[5] Z. Zhang et al., "Jamming ACK Attack to Wireless Networks and a Mitigation Approach," IEEE Global Telecommunications Conf., Nov. 30 - Dec. 4 2008, pp. 4966-4970.
[6] El Hajj Shehadeh et al. "Random backoff control to thwart malicious behavior in WLANs", IEEE Workshop on Local and Metropolitan Area Networks, Apr. 10-12 2013.
[7] J. Guo et al., "A New Trust Management Framework for Detecting Malicious and Selfish Behaviour for Mobile Ad Hoc Networks", IEEE Int. Conf. on Trust, Security and Privacy in Computing and Communications, Changsha, China, 16-18 Nov. 2011, pp.142,149.
[8] WARP: WIRELESS OPEN-ACCESS RESEARCH PLATFORM, [Online].Available: http://warp.rice.edu/.