CROSS-LAYER DESIGNS FOR MITIGATING RANGE ATTACKS IN AD HOC NETWORKS

Jarmo V. E. Mölsä
Communications Laboratory
Helsinki University of Technology
P.O. Box 3000, FI-02015 HUT, Finland
email: jarmo.molsa@tkk.fi

ABSTRACT
This paper analyzes application level performance during a range attack in ad hoc networks. This new denial of service (DoS) attack is based on periodical changes in the transmission range of a wireless node. It is very difficult to detect this attack, because it is perfectly normal for the transmission range to vary. The main contribution of this paper is to present two cross-layer designs for mitigating the range attack. The research methodology is based on using the ns-2 network simulator. Nodes are downloading web pages from a server node in a small simulated ad hoc network, and the complete transmission time for each download is measured. According to the simulation results, the proposed mitigation mechanisms can even double the amount of transmissions fulfilling a specific requirement for maximum delay during a range attack.

KEY WORDS
Wireless networks, Ad hoc routing, Denial of service, Cross-layer design.

1 Introduction

Routing in ad hoc networks is significantly different than in ordinary wired networks [1]. The wireless medium makes it possible to transmit a message to any node within the transmission range of a sender, and dynamic network topologies place heavy requirements on the convergence characteristics of routing protocols. As the name implies, ad hoc networks can be used to construct temporary networks without administrative intervention or any specific infrastructure devices. Correct behavior of all nodes is required, but as expressed in [2], cooperation is assumed but not enforced in mobile ad hoc networks. Even a single malicious node can thus harm routing in a whole ad hoc network. This makes ad hoc routing an attractive target for denial of service (DoS) attacks.

The range attack described in this paper is a new DoS attack against ad hoc routing. It is based on changing periodically the transmission range of a wireless node. There is no need to compromise any node, because an attacker only has to get close enough to the antenna of a node to be used for the range attack. The goal of this attack is to cause frequent topology changes.

This paper analyzes how the range attack affects application level delays when end-users are downloading web pages from a server node. The research methodology is based on using the ns-2 network simulator to measure delays of web page downloads in a small mobile ad hoc network. These delays are compared when the following three ad hoc routing protocols are used: The Ad hoc On-demand Distance-Vector (AODV) [3], the Destination Sequenced Distance-Vector (DSDV) [4], and the Dynamic Source Routing (DSR) [5] protocols. The goal of this analysis is to find out, how vulnerable these ad hoc routing protocols are against the range attack.

The main contribution of this paper is to present two cross-layer designs for mitigating the range attack. A cross-layer design introduces protocol layer interdependencies to optimize overall network performance [6]. Traditionally protocol layers are strictly separated and cannot share network status information between layers.

The first cross-layer design presented in this paper allows the ad hoc routing protocol layer to share information with the medium access control (MAC) protocol layer. The goal here is to assure that unidirectional links are not used by an ad hoc routing protocol, if the underlying MAC layer requires bidirectional links, such as with the IEEE 802.11 MAC protocol. In the simulations both DSDV and DSR encountered transmission breaks when they tried to use unidirectional links caused by the range attack.

The second cross-layer design presented here allows the application layer and the transport layer to share information. If an application is sending time-sensitive information which should be transmitted within a specific maximum delay, an application should wait for the previous transmission to be completely acknowledged before sending a next message.

According to the simulation results, the usage of the presented cross-layer designs can even double the amount of web page transmissions fulfilling a delay requirement during a range attack.

The structure of this paper is the following. First the range attack is specified in more detail. Next, the proposed cross-layer designs are described. The following section explains details about the simulated ad hoc network. After that the simulation results are given. Some related work is also listed. Finally, the conclusions are given.

2 The Range Attack

A range attack is carried out by getting very close to an internal node of an ad hoc network. The node does not need to be compromised. This attack can be carried out by a malicious insider or by an outsider having access to the networked area.

Two kinds of range attacks are studied in this paper: attenuating and amplifying range attacks. In both cases the transmission range of a node is changed periodically between the default range and an attack range. The range attack length defines the time how often the range is changed.

2.1 The Attenuating Range Attack

In an attenuating range attack the transmission range of the attacking node is decreased, for example, by shielding the antenna with some material which prevents or degrades the propagation of the radio signal.

Isolating an internal node periodically from an ad hoc network will cause regular breaks in multi-hop connections traversing through the attacking node. Serious penalties on application level performance will result.

Attenuating range attacks have been analyzed in [7] which compared the resilience of ad hoc routing protocols against this attack and found out that this resilience depends on the requirements of the major applications. This paper continues the work presented in [7].

2.2 The Amplifying Range Attack

In an amplifying range attack the transmission range of the attacking node is increased, for example, by installing an attachment which modifies an omni-directional antenna into a directional antenna. This kind of range extension has been reported, for example, for Bluetooth interfaces, where the default maximum range of 100 meters for a bidirectional connection can be extended up to 500 meters by using a high gain directional antenna [8]. In this paper, however, it is expected that the range extension works only for the transmitted signal, not for the received signal. An extended range can thus only be used for unidirectional links to transmit information to very remote hosts.

The amplifying range attack against ad hoc routing has not been described or analyzed in the literature.

The amplifying range attack should theoretically not cause any transmission breaks, because no links disappear. Some new links are created, but they are all unidirectional. Of course, throughput can decrease due to added interference from amplified radio signals.

In practice, however, the amplifying range attack can cause severe problems in connectivity, especially in ad hoc networks based on MAC protocols requiring bidirectional links for all unicast messages. An example of this kind of a MAC protocol is the IEEE 802.11, in which all unicast transmissions are preceded with the Distributed Coordinated Function (DCF) consisting of request-to-send and clear-to-send signals. All unicast messages are also acknowledged at the MAC level. Only bidirectional links must thus be selected by an ad hoc routing protocol in networks using this kind of a MAC.

Ad hoc routing entities surrounding the attacking node receive route management messages from the attacking node through a unidirectional link. If a routing entity believes that a shorter route is available through the attacking node, it will accept a unidirectional link which cannot forward any messages towards the attacking node. As a result, forwarding of all messages to this link will fail. Depending on the ad hoc routing protocol and the MAC layer, the problem can persist even for the complete duration of the amplified range.

The reason for this vulnerability against the amplifying range attack seems to be the implicit trust for all links being bidirectional. The usage of a bidirectional MAC can increase this false belief. It should be emphasized that all implementations of ad hoc routing protocols do not necessarily have this vulnerability against the amplifying range attack. To prevent this vulnerability, an ad hoc routing protocol should make its own checks for bidirectionality, especially over MAC protocols, such as the IEEE 802.11.

3 Two Cross-Layer Designs for Mitigating the Range Attack

It is very difficult to detect the range attack, because it is perfectly normal for the transmission range to vary due to terrain, moving obstacles, weather etc. For this reason new mitigation mechanisms are needed. This section describes two cross-layer designs for this purpose.

3.1 Routing Level Cross-Layer Design

The routing layer and the MAC layer have at least the following overlapping features regarding the range attack:

- Both layers may have different requirements for bidirectionality.
- Both layers can implement tests for bidirectionality.
- Both layers can implement acknowledgement of transmitted messages.
- Both layers can detect disconnected links

All these features should be coordinated, and a cross-layer design is one possibility for this. These two layers should do not make any false expectations about each other.

The main goal for the proposed routing level cross-layer design is to assure that an ad hoc routing protocol does not accept any unidirectional links, when the ad hoc network is based on a a MAC requiring bidirectional links. This is necessary for mitigating or even preventing the amplifying range attack. When necessary, an ad hoc routing agent should implement itself additional features to prevent

unidirectional links from entering route tables. Also, cooperation between these two layers make it possible to implement acknowledgements and link status detection in the most effective place. Cross-layer designs can make it easier to detect inconsistencies between the properties of these two layers.

3.2 Application Level Cross-Layer Design

When there is no end-to-end path available to a destination (e.g., temporarily due to the range attack), all packets to that destination will be queued until a path is again available. The associated delay can be problematic especially for an application sending time-sensitive data, such as regular updates for some information. The proposed application level cross-layer design will reduce the possibility for a message being delayed excessively.

The main goal of the application level cross-layer design is to prevent an application from sending new time-sensitive data when previous messages have not yet been acknowledged at the TCP level. Otherwise a new message would just remain in a send buffer waiting for a usable end-to-end path, and during this time the message would gradually lose its timeliness and waste network resources unnecessarily. This cross-layer design thus involves sharing TCP acknowledgement status with an application. When transmitting messages relatively infrequently, a sign of an unacknowledged previous message tells about an unavailable end-to-end path.

An alternative to providing TCP acknowledgement status to an application is the use of the Stream Control Transmission Protocol (SCTP) [9] instead of TCP. SCTP can reduce problems from the head-of-line blocking where a lost packet prevents packets with higher sequence numbers from being delivered to an application, even if these packets would have been received correctly.

4 The Simulated Ad Hoc Network

The ns-2.28 network simulator was used to investigate the application level performance during range attacks. Two modifications were made to the basic ns-2.28 simulator: nodes were allowed to have different transmission ranges, and the infinite loop problem of the DSDV was patched.

The structure of the simulated ad hoc network is shown in the Fig. 1. This network consists of six nodes, numbered from 0 to 5. The x- and y-coordinates for a node are indicated in parenthesis below each node. The IEEE 802.11 MAC layer is used in the network. All messages are transmitted with the bandwidth of 1 Mbps.

Nodes 0 to 4 are static. The node 5 is moving vertically along the y-axis back and forth between the points (300,700) and (300,100). At the beginning of a simulation it starts moving downwards with the speed of 3 m/s. At the time of 400 seconds it starts moving upwards. The node 5 initiates a movement every 400 seconds.

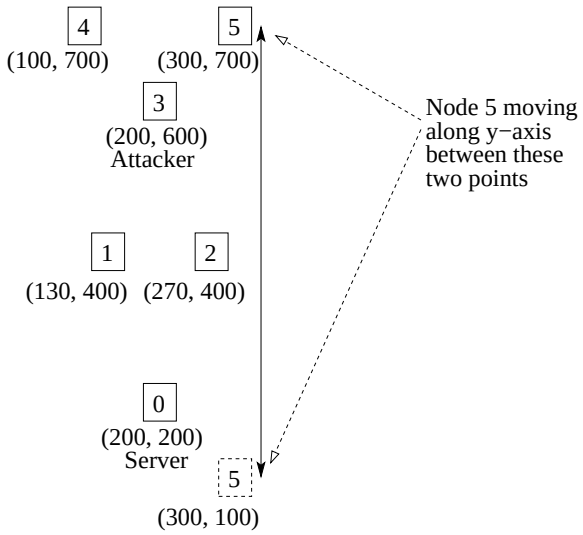

Figure 1. Structure of the simulated ad hoc network.

The node 3 is used for the range attack. The default transmission range for all nodes is 250 meters. In the attenuation range attack the range of the node 3 is reduced periodically to 40 meters. In the amplifying attack this range is periodically increased to 550 meters.

Client nodes are downloading web pages from the server node 0 with an exponentially distributed inter-page time, the average value being 30 seconds. These pages are downloaded automatically over the TCP protocol. Each web page contains 2920 bytes, which results in two full-size TCP segments. It is expected that persistent TCP connections are used, so the three-way handshake is not required for initiating a download. It should be noticed that the downloaded information does not necessarily have to be a web page, because an application is only expected to use TCP for its transmission purposes.

The transmission delay for a download is the complete time to transmit and acknowledge a single web page. This delay is thus the time from the transmission of the first TCP segment to the reception of the acknowledgement of the second TCP segment at the server node 0.

4.1 Simulation Parameters

The length of each simulation was 60 000 seconds. Due to the memory requirements of DSR, every simulation was divided into 15 independent sub-simulations, each of length 4 000 seconds. All sub-simulations used different random number sequences.

Simulations were repeated with the following parameter combinations:

- The routing protocol was AODV, DSDV, or DSR.
- The range attack length (time between range modifications) was 1, 5, 10, 20, 30, 40, 50, or 60 seconds.

5 Simulation Results

In all of the simulations, the transmission delay of a web page was measured only to the node 4. The node 4 suffers the most from the range attack, as it uses frequently a multi-hop connection through the attacking node 3. A multi-hop connection through the mobile node 5 is also possible, depending on the location of the node 5.

All graphs have three different curves. The continuous line represents the AODV protocol, the narrow dotted line represents the DSDV protocol, and the broad dotted line represents the DSR protocol.

5.1 Normal Delay

The Fig. 2 shows the cumulative distribution function (CDF) for the delay of transmissions to the node 4. This figure shows the normal delay distribution when there is no attack. $CDF(x)$ is the probability that the transmission delay is less or equal to x.

As can be seen from the Fig. 2, the complete transmission delay is mostly below 0.1 seconds. In case of the DSDV protocol, approximately 5% of downloads experience a relatively long delay of more than 10 seconds.

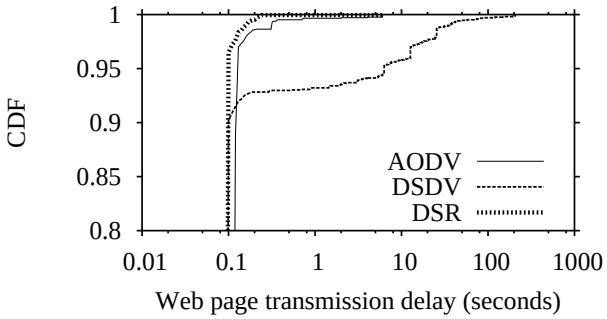

Figure 2. Cumulative Distribution Function for the delay of the node 4 when there is no range attack.

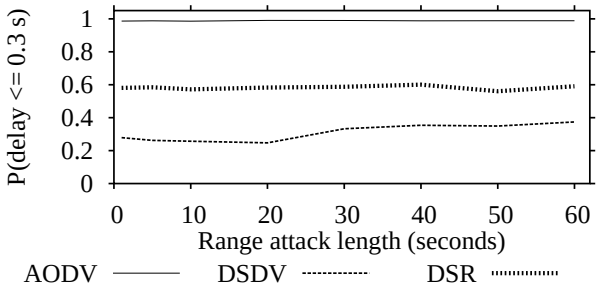

Figure 3. Fraction of transmissions having a delay $\leq$ 0.3 s for the node 4 during an amplifying range attack.

5.2 Delay During an Amplifying Range Attack

The Fig. 3 shows the fraction of transmissions having a delay less or equal to 0.3 seconds during an amplifying range attack. In this case it is expected that the application is transmitting very time-sensitive information. The x-axis in this and all the following figures shows the range attack length.

The amplifying range attack should have no effect on the delay distribution because no links are broken due to this attack. All end-to-end paths remain available regardless of this attack. These results indicate, however, that the implementations of both DSDV and DSR in the ns-2 network simulator are vulnerable to this attack. The AODV protocol is completely insensitive to the amplifying range attack.

The Fig. 4 shows the fraction of transmissions having a delay less or equal to 2 seconds. When comparing figures 3 and 4 we can see the following:

- AODV is completely insensitive to the amplifying range attack.
- DSR tries to use unidirectional links for short periods of time. It takes approximately 0.6–1.3 seconds to recover from this situation. The effect from the amplifying range attack is visible only in the Fig. 3 where applications require a very short delay. If applications tolerate a delay less or equal to 2 seconds, DSR is practically insensitive to this attack.

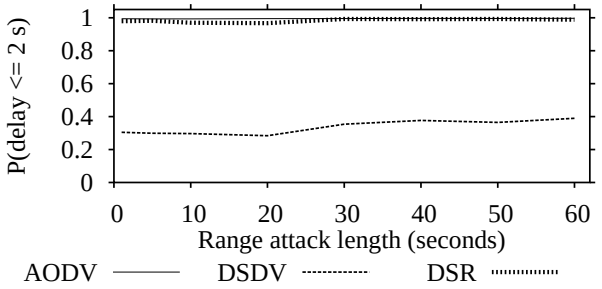

Figure 4. Fraction of transmissions having a delay $\leq$ 2 s for the node 4 during an amplifying range attack.

- The application level performance collapses with DSDV in ns-2.

These results show that ad hoc routing protocols can be vulnerable to the amplifying range attack, even if it should not be possible according to the specifications. This finding in the ns-2 network simulator should be treated as a proof of concept, that the amplifying range attack can cause denial of service. The use of the proposed routing level cross-layer design would prevent this kind of a DoS attack, if unidirectional links are rejected immediately.

5.3 Delay During an Attenuating Range Attack

Figures 5 and 6 show the fraction of transmissions having a delay less or equal to 0.3 or 2 seconds, respectively. An attenuating range attack is carried out from the node 3.

As can be seen from these two figures, long delays result from the attenuating range attack. Approximately 50% of the transmissions fail to have a delay less or equal to 2 seconds. This result is not a surprise, because losing connectivity in an ad hoc network will certainly increase delays.

5.4 Benefits from the Application Level Cross-Layer Design

All previous simulation results were achieved without using the proposed application level cross-layer design. This subsection shows, how delay properties change, when an application waits for the previous transmission to finish completely before initiating a new one. The higher the download frequency, the better the effect from using this design. In the remaining figures the average inter-page time is 5 seconds.

Figures 7 and 8 show the fraction of transmissions having a delay less or equal to 2 seconds when the application level cross-layer design is not included and when it is included, respectively.

When the application level cross-layer design is not used, at least 20–30% of the transmissions fulfill the delay requirement. When the application waits for the previous transmission to be acknowledged, at least 60% of the transmissions fulfill the delay requirement. As a result, the amount of transmissions fulfilling the delay requirement has been doubled by using this cross-layer design between application and transport layers.

5.5 Results When All Client Nodes Moving

All previous simulations include only one moving node. In this second simulation setup all client nodes 1–5 are moving vertically. During movement the y-coordinate of a node changes back and forth between the initial position of the node (see Fig. 1) and the value of 100. All client nodes are moving with the speed of 3 m/s. Nodes 1 and 2 initiate a movement every 150 seconds. Nodes 3–5 initiate a movement every 400 seconds.

In this scenario the higher mobility increases the probability of long delays for the node 4. Approximately 80% of downloads have a delay less or equal to 10 seconds when there is no attack. The node 4 is here more tolerant against the range attack because it has a single-hop connection to the server node part of the time. The application level cross-layer design improved the performance here by 20–30 percent units, when average inter-page time was 5 s.

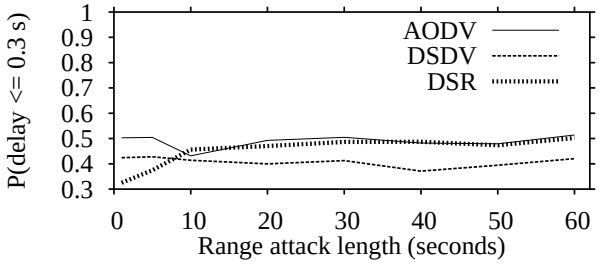

Figure 5. Fraction of transmissions having a delay $\leq$ 0.3 s for the node 4 during an attenuating range attack.

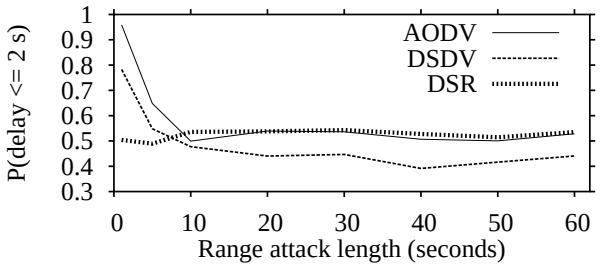

Figure 6. Fraction of transmissions having a delay $\leq$ 2 s for the node 4 during an attenuating range attack.

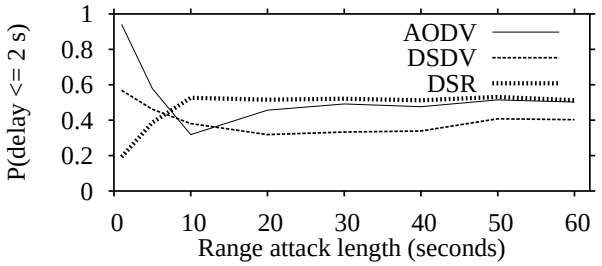

Figure 7. Fraction of transmissions having a delay $\leq$ 2 s for the node 4 during an attenuating range attack. Average inter-page time is 5 s. There is no interaction between application and transport layers.

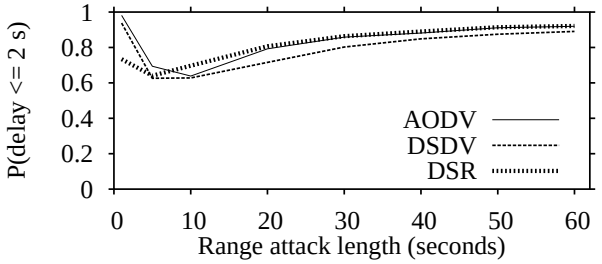

Figure 8. Fraction of transmissions having a delay $\leq$ 2 s for the node 4 during an attenuating range attack. Average inter-page time is 5 s. Application waits before previous transmission completely finished.

6 Related Work

Studies about DoS attacks in ad hoc networks have mostly concentrated on the misuse of routing protocols, such as injecting false routing messages [10, 11, 2]. Few papers have studied other types of DoS attacks in ad hoc networks. For example, the jelly fish attack forces TCP flows to have almost zero throughput by simply reordering, dropping, or causing variable delay to forwarded TCP segments [12]. Mitigation of route request floods in ad hoc networks was studied in [13].

Several cross-layer designs have been proposed for ad hoc networks. Most of these designs concentrate on improving throughput through cross-layer power control, such as in [14]. An architecture for a cross-layer design covering all possible layers is described in [6].

Directional antennas have the possibility of improving the performance significantly when compared to omnidirectional antennas. Directional antennas can also be used to overcome some DoS attacks, such as jamming. A system solution for using directional antennas in ad hoc networks is presented in [15].

7 Conclusion

This paper analyzed application level performance in ad hoc networks during range attacks. These attacks are based on modifying the transmission range of a wireless node periodically.

Simulations with the ns-2 network simulator revealed that ad hoc routing protocols can accidentally accept unidirectional links when an attacking node increases its transmission range. Only bidirectional links should be used over the IEEE 802.11 MAC protocol. This vulnerability is clearly visible with DSDV which suffers from a large degradation in performance during an amplifying range attack. The DSR protocol also tried to use unidirectional links for 0.6–1.3 seconds, which increased the average delay of application level message transmissions. A cross-layer design was proposed in this paper to enforce checks for this important issue and prevent this attack completely. This design makes it possible for the ad hoc routing and MAC protocol layers to share information about each other to limit the possibility for a mismatch.

Another cross-layer design was proposed here for the application and transport layers. If an application is sending time-sensitive information which loses its meaning gradually in time, a new message should only be sent when previous messages have been completely acknowledged at the transport layer. This design was tested over the TCP protocol. By using this cross-layer design it was possible to even double the amount of transmissions fulfilling a delay requirement of 2 seconds.

According to these simulation results, cross-layer designs can effectively mitigate DoS attacks against ad hoc routing.

Acknowledgements

This work was supported by the Finnish Defence Forces. The author would like to thank both the members of the project group and the anonymous reviewers for their helpful comments in improving this paper.

References

[1] C. E. Perkins, Ed., *Ad Hoc Networking* (Upper Saddle River, New Jersey, USA: Addison-Wesley, 2001).

[2] H. Yang, H. Luo, F. Ye, S. Lu, & L. Zhang, Security in mobile ad hoc networks: Challenges and solutions, *IEEE Wireless Communications, 11*(1), 2004, 38–47.

[3] C. Perkins, E. Belding-Royer, & S. Das, Ad hoc on-demand distance vector (AODV) routing, The Internet Society, RFC 3561, 2003.

[4] C. E. Perkins, & P. Bhagwat, Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers, *Proc. SIGCOMM '94 Conference on Communications Architectures, Protocols and Applications*, London, England, UK, 1994, 234–244.

[5] D. Johnson, D. Maltz, & Y.-C. Hu, The dynamic source routing protocol for mobile ad hoc networks (DSR), 2004, internet draft draft-ietf-manet-dsr-10.txt, work in progress.

[6] M. Conti, G. Maselli, G. Turi, & S. Giordano, Cross-layering in mobile ad hoc network design, *IEEE Computer, 37*(2), 2004, 48–51.

[7] J. V. E. Mölsä, Increasing the DoS attack resiliency in military ad hoc networks, *Proc. IEEE MILCOM*, Atlantic City, New Jersey, USA, 2005.

[8] M. Bialoglowy, Bluetooth security review, part 2, Security Focus, Tech. Rep., 2005.

[9] R. R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. J. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, & V. Paxson, Stream control transmission protocol, The Internet Society, RFC 2960, 2000.

[10] W. Wang, Y. Lu, & B. K. Bhargava, On security study of two distance vector routing protocols for mobile ad hoc networks, *Proceedings of the IEEE PerCom*, Fort Worth, Texas, USA, 2003, 179–186.

[11] Y.-C. Hu, A. Perrig, & D. B. Johnson, Rushing attacks and defense in wireless ad hoc network routing protocols, *Proc. ACM WiSe*, San Diego, California, USA, 2003, 30–40.

[12] I. Aad, J.-P. Hubaux, & E. W. Knightly, Denial of service resilience in ad hoc networks, *Proc. ACM MobiCom*, Philadelphia, Pennsylvania, USA, 2004, 202–215.

[13] S. Desilva, & R. V. Boppana, Mitigating malicious control packet floods in ad hoc networks, *Proc. IEEE Wireless Communications and Networking Conference*, New Orleans, Louisiana, USA, 2005.

[14] V. Kawadia, & P. R. Kumar, Principles and protocols for power control in wireless ad hoc networks, *IEEE Journal on Selected Areas in Communications, 23*(1), 2005, 76–88.

[15] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, & S. Polit, Ad hoc networking with directional antennas: A complete system solution, *IEEE Journal on Selected Areas in Communications, 23*(3), 2005, 496–506.