The Effects of Flooding Attacks on Time-Critical Communications in the Smart Grid

Qinghua Li*, Chase Ross*, Jing Yang[†], Jia Di*, Juan Carlos Balda[†], and H. Alan Mantooth[†]

* Department of Computer Science and Computer Engineering, University of Arkansas

[†] Department of Electrical Engineering, University of Arkansas

Emails: qinghual@uark.edu, cmross@email.uark.edu, {jingyang, jdi, jbalda, mantooth}@uark.edu

Abstract—Many smart grid communications are delay sensitive and have very strict timing requirements for message deliveries. For example, trip protection messages must be delivered to the destination within 3 ms according to IEC 61850. Such time-critical communications are vulnerable to flooding attacks which attempt to increase message delivery delay through congesting the network channel and exhausting the computation resources of the communicating nodes. However, there is a lack of understanding on how much flooding attacks affect message delivery delays. In this paper, we conduct experimental studies to investigate how flooding attacks affect message delivery delays for time-critical communications in smart grid. Our experiments are based on both wireless networks in a lab and wired networks in a real, industry-standard electric power facility. Experimental results show that even low-rate flooding attacks can significantly increase the message delivery delays, especially when wireless networks are used.

I. INTRODUCTION

Many smart grid communications (especially in substation automation systems) are delay sensitive and have very strict timing requirements for message deliveries. For example, messages for trip protection have the delay constraint of 3 ms and station-wide interlocking messages should be delivered to relays in 10 ms according to IEC 61850 [1]. For such time-critical messages, their delay requirement must be met to prevent catastrophic consequences for the power infrastructure.

However, flooding attack can increase the delays of time-critical messages. Flooding attack is a type of denial-of-service (DoS) attack, in which the attacker floods a lot of messages to a target network or a target machine. Flooding attack affects the delivery delay of time-critical messages mainly in two ways. First, flooded messages consume network resources (e.g., bandwidth, network channel, and buffers in switch). Second, when the target is the sender and receiver of time-critical messages, flooded messages consume the sender's and the receiver's local resources (e.g., CPU cycles). For instance, if a lot of messages are flooded to a circuit breaker to congest its communication channel and exhaust its computation resources, it may not be able to receive legitimate open/close commands from the protective relay in a timely manner during abnormal events (e.g., fault). Thus, it is important to understand *how much effect flooding attack has on message delivery delays*.

A lot of work has been done to study flooding attacks (e.g., [2]–[4]), but they mainly focus on how flooding attack affects message delivery ratio. Some studies discussed the effect on message delay a little as a small part of a big evaluation

framework, but to the best of our knowledge there is still no systematic study about flooding attack's effect on message delivery delays. Especially, there is no such study in the context of smart grid considering the unique time constraint of power grid communications. As a result, many important questions remain open. For example, how much traffic must an attacker flood to make the message delay higher than its limit? What type of flooding (e.g., flooding to network or to target host) should the attacker launch to achieve this goal?

In this paper, we conduct experimental study of how flooding attack affects message delivery delays for time-critical communications. We carefully design experiments to unveil flooding attack's effect on various components of message delivery delay, develop flooding attack tools, and conduct systematic experiments in wireless and wired networks. This paper makes the following contributions:

- This is the first systematic, experimental study of how flooding attack affects message delivery delays in the context of smart grid. It considers network-level flooding and application-level flooding attacks, and it evaluates TCP, UDP, and Raw MAC (defined later) protocols on both wireless and wired networks.
- To the best of our knowledge, it provides the first experimental results on a real, industry-standard power facility - the National Center for Reliable Electric Power Transmission (NCREPT) [5] at the University of Arkansas.
- The obtained understandings will help utilities evaluate the risk of their power infrastructure being endangered by flooding attacks, and provide guidance for correctly configuring their countermeasures against flooding attacks. In particular, this paper identifies the threshold flooding rate beyond which the delay will be higher than the delay constraint of time-critical communications. This can help set the right threshold for filtering flooded traffic.

The rest of this paper is organized as follows. Section II introduces background knowledge on message delivery delay. Section III presents our experiment design. Section IV presents experiment results. Section V presents related work. The last section concludes the paper.

II. PRELIMINARIES ON MESSAGE DELIVERY DELAY

Figure 1 shows how a message is sent from the sender and delivered to the receiver. Normally, a message generated by the sender application is passed down the protocol stack

Fig. 1. Illustration of how a message is delivered from the sender to receiver.

layer-by-layer to the Medium Access Control (MAC)/physical layer (which runs in the network interface card). Then the message is transmitted to the network interface at the receiver. At the receiver, it is passed on the application layer again in a layer-by-layer style. In this process, each layer incurs some processing time which contributes to the overall message delivery delay (whose definition will follow). To reduce the delay and meet the requirement of time-critical messages (e.g., GOOSE), IEC 61850 directly passes a message from the application layer to the MAC/physical layer at the sender and vice versa at the receiver. For convenience, we denote such a protocol as Raw MAC in this paper. These two modes are illustrated in Figure 1 as “regular path” and “direct path”.

Message delivery delay is defined as the elapsed time from when a message is sent out by the sender application to when the message is received by the receiver application. It contains several main parts. 1) Processing time at the sender’s network protocol stack. This denotes the time needed for the message to be conveyed from the application layer to the network interface. 2) Network delay. After the message reaches the sender’s network interface, since there might be some other messages queued for transmission, the message must wait until the messages before it are transmitted. When the message is at the head of the sending queue, some time is still needed for the message to be received by the receiver’s network interface. 3) Processing time at the receiver’s network protocol stack. This is time needed for the message to be delivered from the network interface to the application layer. When there are many messages waiting to be delivered, the message may be buffered for some time before being delivered.

III. EXPERIMENT DESIGN

A. Measurement of One-Way Message Delivery Delay

We aim to measure the one-way message delivery delays between the sender application and the receiver application. There are many tools available for measuring the round-trip time (e.g., [6]), but they usually work in the network layer and hence do not meet our needs. Also, they are not designed for time-critical communications, and it is not sure if these tools will induce extra delays. Thus, we design our own algorithms to measure the delays of TCP, UDP, and Raw MAC protocols.

For TCP, we run two TCP streams between two endpoints Alice and Bob. For Stream 1, Alice is the client and Bob is the server; for Stream 2, Alice is the server and Bob is the client. The one-way delay is measured as follows. Alice records its current local time as the sending time t_1 , initiates a

TCP connection with Bob, and sends a short message (with 32 bytes of payload) to Bob. Bob receives this message, initiates another TCP connection with Alice, and sends the same message to Alice. Alice receives the message, and records its current local time as the receiving time t_2 . Then the one-way delay is calculated as $\frac{t_2 - t_1}{2}$.

For UDP, the process is similar to TCP. Alice records its current local time as the sending time t_1 , and sends a short UDP message (with 32 bytes of payload) to Bob. Bob receives this message, and sends the same message to Alice via another UDP socket and port. Alice receives the message, and records its current local time as the receiving time t_2 . Then the one-way delay is calculated as $\frac{t_2 - t_1}{2}$. For Raw MAC, the process is exactly the same as UDP.

B. Flooding Attacks

Section II shows that message delivery delay consists of three portions, processing time at sender, network delay, and processing time at receiver. To study the effect of flooding attacks on each of these portions, we consider network-layer flooding and application-layer flooding attacks. In network-layer flooding, the flooded packets will not go to the application layer of the sender or receiver. To the contrast, in application-layer flooding, packets will be flooded to an application-layer service run at the sender or receiver and will be processed at the application layer.

We develop the following types of flooding attacks:

- **Flooding to a third node** (network-layer). The attacker floods UDP packets to a third node in the network, which is not the sender or receiver of time-critical communications. For the sender and receiver, the flooded messages are only background traffic; however, such traffic still contends with legitimate messages for network channel.
- **Broadcast flooding** (network-layer). The attacker broadcasts UDP packets to the whole network.
- **Unsolicited flooding to sender** (network-layer). The attacker floods UDP packets to the sender. Since the sender has no service accepting such packets, these packets will not go to the application layer at the sender.
- **Unsolicited flooding to receiver** (network-layer). This is similar to the previous case, but this time the attacker floods UDP packets to the receiver.
- **Solicited flooding to sender** (application-layer). The attacker floods UDP packets to the sender. The sender runs certain service at the arriving port and hence these packets will be processed in some way. In our implementation, the sender simply prints out the received packet.
- **Solicited flooding to receiver** (application-layer). This is similar to the previous case, but this time the attacker floods UDP packets to the receiver.

C. Deployment Scenarios

In recent years, many studies have proposed to use wireless networks to deliver messages in the power grid infrastructure [7]–[9]. Although modern wireless networks are efficient and easy to configure, their vulnerability to flooding attacks is

Fig. 2. Network-layer flooding against Raw MAC communications.

not well understood yet when they are used in time-critical applications such as substation automation systems. To address this problem and evaluate the risk of using wireless networks for time-critical communications, we deploy the experiments in a Wireless Local Area Network (WLAN).

On the other hand, many electric substations use switched Ethernet as the communications network. To evaluate their vulnerability to flooding attacks, we also deploy the experiments on a switched Ethernet.

IV. IMPLEMENTATION AND EXPERIMENT RESULTS

A. Implementation

We implement a flooding attack tool on Linux platforms with the C/C++ programming language. The tool implements all the flooding attacks described in Section III-B. We also implement all the algorithms described in Section III-A which measure the one-way delivery delay of TCP, UDP, and Raw MAC protocols. Our testbed consists of four laptop computers. The sender of time-critical messages (denoted by Alice) is implemented on a Dell Latitude-E6440 laptop with an Intel Core i7-4610M CPU (3 GHz). The receiver of time-critical messages (denoted by Bob) is implemented on a Dell Precision-M6500 laptop with an Intel Core i7 CPU Q 820 (1.73GHz x 8). The attacker is implemented on an HP-G60 with a Pentium Dual-Core CPU T4200 (2 GHz x 2). The third node (denoted by Carl) does not run any special code, and we use a Dell Vostro laptop to play its role.

B. Experiments on a Wireless LAN

In this section, we conduct experiments and evaluate how flooding attacks affect message delivery delays in a WLAN.

1) *Experimental Setting*: We build a WLAN using a Linksys E1000 wireless router. The four laptops (Alice, Bob, Attacker, and Carl) are connected to it. The size of flooded packets is 1000B.

2) *Results for Raw MAC*: In this section, we evaluate how flooding attacks affect the delays of Raw MAC.

Network-layer flooding We first study the effect of network-layer flooding attacks. Specifically, *flooding to a third node* and *broadcast flooding* are considered here. The first attack floods UDP packets to the third node Carl, and the second attacks floods UDP packets to the whole network using the broadcast address. Figure 2 shows the results. In this figure, X-axis is the flooding rate in packets/second, and Y-axis is the average message delivery delay with 95% confidence interval.

Fig. 3. Application-layer flooding against Raw MAC communications.

When flooding to a third node, the message delivery delay increases as the flooding rate increases. This is because WLAN has a single collision domain, and flooded packets compete the shared wireless channel with the packets sent between Alice and Bob. The higher flooding rate, the more contention, and the higher message delivery delay.

When flooding to the whole network, the message delivery delay also increases as the flooding rate increases for similar reasons. However, we can see that at the same flooding rate the delay in broadcast flooding is much higher than the delay in flooding to a single node. For example, the difference is of two magnitudes of order when the flooding rate is 512 packets/second. This is because broadcast messages consume more network bandwidth and channel resources than unicast.

It can also be seen that, when there is no attack, the delay is only 2.6ms, which is less than the 3ms limit set by IEC 61850. However, under both attacks, the delay can easily get to exceed that limit. Thus, flooding attack is destructive to time-critical communications.

Application-layer flooding Then we study the effect of application-layer flooding attacks, i.e., *solicited flooding to sender* and *solicited flooding to receiver*. Both flood UDP packets. The results are shown in Figure 3. For both attacks, the delay increases as the flooding rate increases. There are two reasons for this trend. One reason is that the flooded packets contend for network channels with legitimate packets between Alice and Bob. The other reason is that Alice and Bob need to process the flooded packets when they receive them. This consumes their CPU cycles and hence can delay the processing of legitimate packets.

For comparison, the result of *flooding to third node* is also plotted in Figure 3. We can see that, when unicast is used, application-layer flooding induces longer message delivery delays than network-layer flooding, especially when the flooding rate is high. The reason is that in application-layer flooding attacks, flooded packets also contend for local resources at the sender and receiver.

It can be seen that, when flooding packets to the receiver, the delay is longer than when flooding packets to the sender, although the caused network contention is the same in both cases. The reason is that when packets are flooded to the receiver, the flooded packets will contend for the delivery queue between the receiver's network interface and application layer. This will more delay the time when a legitimate packet

Fig. 4. Critical flooding rate under broadcast flooding against Raw MAC communications.

is delivered to the application layer.

Critical flooding rate In all the above cases, the delay under flooding attack can exceed the 3ms limit. It is interesting to see at which rate the delay will go beyond that limit. For *flooding to third node*, *solicited flooding to sender*, and *solicited flooding to receiver*, the transition happens at rate 128 packets/second, which can be found from Figure 2 and 3. To identify the transition rate for *broadcast flooding*, we run a set of experiments at lower flooding rates. Figure 4 shows the results. It can be seen that the transition rate is 4 packets/second. This rate is very low, which indicates that the attacker can slowly flood the network to make the delay of time-critical messages too long to be useful. *Such low rate might render traditional broadcast rate limit (which is usually larger than 4 packets/second) useless.*

3) *Results for TCP and UDP:* In this section, we evaluate how network-layer and application-layer flooding attacks affect the delays of TCP and UDP communications. Similar to the Raw MAC case, we evaluate *flooding to a third node* and *broadcast flooding* as the two network-layer flooding attacks and *solicited flooding to sender* and *solicited flooding to receiver* as the two application-layer flooding attacks. The results for TCP communications are shown in Figure 5. The results for UDP communications are shown in Figure 6. The results have similar trends as the Raw MAC case.

4) *Summary:* The above results show that time-critical communications in WLAN are vulnerable to flooding attacks. For network-layer flooding attacks, broadcast flooding has much more severe effect than flooding to third node at the same flooding rate. For application-layer flooding attacks, solicited flooding to receiver tend to have more severe effect than solicited flooding to sender at the same flooding rate. Additionally, application-layer flooding tends to induce longer delays than network-layer flooding at the same flooding rate.

C. Experiments on a Real Power System Network

In this section, we conduct experiments and evaluate how flooding attacks affect message delivery delays in a switched Ethernet which is used by a real, industry-standard power facility - NCREPT.

1) *NCREPT:* The NCREPT center has distribution-level (15 kV/300 A) test facilities which house regen drives, circuit breakers, protection relays, transformers, controls, data acquisition units, etc. Key to providing the flexibility and

Fig. 5. Flooding attacks against TCP communications.

Fig. 6. Flooding attacks against UDP communications.

reconfigurability required for this facility are the 6 MVA regenerative drives, which can generate various resistive, inductive, and capacitive loads. With them one can generate a range of loads present on the grid consistent from factories using heavy machinery for manufacturing to residential homes. Various industry standard communication protocols are used to monitor and control all of the aforementioned equipment, such as TCP/IP, Modbus, and Incom which are commonly found in power substations and industry.

The network architecture of NCREPT is illustrated in Figure 7(a). The core of this network is a switched Ethernet LAN which contains two RuggedCom RS900 industrial grade Ethernet switches. The two switches are connected with fiber optics. Circuit breakers and protective relays are connected to Switch 2 via two EMINT Ethernet to Incom network translators. PLC controllers and other devices are also connected to Switch 2. System monitor is connected to Switch 1.

2) *Experimental Setting:* Our experimental setting is also shown in Figure 7(a). The communications between protective relays and circuit breakers are usually time-critical. To reflect the communications between them, we connect the sender and receiver (Alice and Bob) of time-critical messages in our experiments to Switch 2. We also connect Attacker to Switch 2. The third node Carl is connected to Switch 1 but it is not shown in this figure for simplicity. A snapshot of our testbed in NCREPT is shown in Figure 7(b). The fourth laptop is not included in this picture since it is located in another room.

3) *Results:* We first study the effect of network-layer flooding attacks. Since there is no collision domain in a switched Ethernet, *flooding to a third node* does not affect delay much. Instead, we consider *unsolicited flooding to sender* and *unsolicited flooding to receiver* here. Figure 8 shows the results. We can see that as the flooding rate increases, the delay

Fig. 7. Our experimental setting in the NCREPT center.

Fig. 8. Network-layer flooding attacks against Raw MAC communications in NCREPT.

increases under both attacks. The reason is as follows. When packets are flooded to the sender (receiver), flooded packets can still contend with legitimate packets for the link between the switch and the sender (receiver), and induce extra delays.

We then study the effect of application-layer flooding attacks. Here, we consider *solicited flooding to sender* and *solicited flooding to receiver*. Figure 9 shows the results. Again, when the flooding rate increases, the delay significantly increases under both attacks. In these attacks, flooded packets not only contend with legitimate packets for the link between the switch and the sender (receiver), but also for CPU cycles at the sender and receiver. Due to the extra resources contended for, application-layer flooding attacks induces longer delays than network-layer flooding attacks. For instance, at rate 11637 packets/second, the delay under application-layer attacks is larger than 1ms, which is twice of the delay under network-layer attacks.

V. RELATED WORK

Time-critical communications in smart grid have been studied by recent work [10]–[12]. Lu et al. [10], [11] study how jamming attacks affect time-critical communications, but they do not study flooding attacks, and do not provide systematic experimental results on how attacks affect delays. Georg et al [12] propose a modeling approach for evaluating high performance and real-time capability of communication technologies, but they do not study flooding attacks either.

Fig. 9. Application-layer flooding attacks against Raw MAC communications in NCREPT.

VI. CONCLUSIONS

In this paper, we studied the effect of flooding attacks on message delivery delays for time-critical communications in smart grid. We designed various network-layer and application-layer flooding attacks, and deployed these attacks in a WLAN and a switched Ethernet LAN in a real power facility. Experimental results show that flooding attacks can significantly increase the delay of time-critical communications, especially when WLAN is used.

ACKNOWLEDGMENT

We thank Chris Farnell for helping us set up the experiments in the NCREPT center.

REFERENCES

- [1] I. S. 61850, "Communication networks and systems in substations," 2003.
- [2] A. Afanasyev, P. Mahadevan, I. Moiseenko, E. Uzun, and L. Zhang, "Interest flooding attack and countermeasures in named data networking," in *IFIP Networking Conference*, 2013, May 2013, pp. 1–9.
- [3] P. Gasti, G. Tsodik, E. Uzun, and L. Zhang, "Dos and ddos in named data networking," in *Computer Communications and Networks (ICCCN)*, 2013 22nd International Conference on, July 2013, pp. 1–7.
- [4] M. Jensen, N. Gruschka, and N. Luttenberger, "The impact of flooding attacks on network-based services," in *Availability, Reliability and Security*, 2008. *ARES 08. Third International Conference on*, March 2008, pp. 509–513.
- [5] N. C. for Reliable Electric Power Transmission, "http://ncrept.uark.edu/."
- [6] Netperf, "http://www.netperf.org/netperf/."
- [7] M. Tanaka, D. Umehara, M. Morikura, N. Otsuki, and T. Sugiyama, "New throughput analysis of long-distance ieee 802.11 wireless communication system for smart grid," in *Smart Grid Communications (SmartGridComm)*, 2011 IEEE International Conference on, Oct 2011, pp. 90–95.
- [8] B. Akyol, H. Kirkham, S. Clements, and M. Hadley, "A survey of wireless communications for the electric power system," 2010.
- [9] P. Kanabar, M. Kanabar, W. El-Khattam, T. Sidhu, and A. Shami, "Evaluation of communication technologies for iec 61850 based distribution automation system with distributed energy resources," in *Power Energy Society General Meeting, 2009. PES '09. IEEE*, July 2009, pp. 1–8.
- [10] Z. Lu, W. Wang, and C. Wang, "Modeling, evaluation and detection of jamming attacks in time-critical wireless applications," *Mobile Computing, IEEE Transactions on*, vol. 13, no. 8, pp. 1746–1759, Aug 2014.
- [11] —, "From jammer to gambler: Modeling and detection of jamming attacks against time-critical traffic," in *INFOCOM, 2011 Proceedings IEEE*, April 2011, pp. 1871–1879.
- [12] H. Georg, N. Dorsch, M. Putzke, and C. Wietfeld, "Performance evaluation of time-critical communication networks for smart grids based on iec 61850," in *INFOCOM, 2013 Proceedings IEEE*, April 2013, pp. 3417–3422.