Analysis of the Effects of Distributed Denial-of-Service Attacks on MPLS Networks

Béla Genge[1] and Christos Siaterlis

Institute for the Protection and Security of the Citizen, Joint Research Centre, Via Enrico Fermi 2749, Ispra (Varese) 21027, Italy

Abstract

Modern critical infrastructures such as the power grid are frequently targeted by distributed denial-of-service (DDoS) attacks. Unlike traditional information and communications systems, where the effects of DDoS attacks are mostly limited to the cyber realm, disruptive attacks on critical infrastructure assets can result in the loss of vital services such as transportation and health care. This paper evaluates the effect of disruptive DDoS attacks on multiprotocol label switching (MPLS) networks that provide communications services to many large-scale critical infrastructure assets. The experimental results provide insights into architectural configurations that can increase network resilience without the need to incorporate additional hardware and software.

Keywords
Distributed Denial-of-Service Attacks; MPLS Networks; Resilience
Submitted: December 31, 2012
Accepted: April 10, 2013

1. Introduction

Over the last few years there has been a considerable increase in the scope and impact of distributed denial-of-service (DDoS) attacks. These attacks

[1] Corresponding author: Béla Genge (bela.genge@jrc.ec.europa.eu)

flood targeted systems with packets from thousands of different sources, disrupting communications and services and bringing down even well-defended targets. An example is the recent attack on Spamhaus [10], a massive 300 Gbps packet flood that saturated the company's Internet connections and blocked access to its web pages. The attack was rated as possibly the largest DDoS attack in history. It clearly illustrates the effectiveness of DDoS attacks that employ legitimate services such as DNS to achieve their ends.

Unfortunately, modern critical infrastructures such as power grids, oil and gas pipelines, and water supply systems are constantly exposed to DDoS attacks. A November 2010 study conducted by McAfee [1] involving 200 industry executives in fourteen countries revealed that more than 80% of critical infrastructure installations faced DDoS attacks that year. Part of the problem is that security and resilience measures are not made compulsory through policy or regulation, and telecommunications operators are often not aware of the severe risks imposed by the lack of security mechanisms. The situation is exacerbated by the fact that security mechanisms are commonly misconfigured [5, 15], potentially rendering the entire security posture ineffective.

The research community has proposed several approaches for designing resilient network topologies. Some of these approaches employ network traffic models that limit input traffic flow to ingress routers based on the available bandwidth [9, 14]. However, disruptive DDoS attacks take the hardware and software to their limits, resulting in system states that are difficult to analyze using existing approaches.

In an attempt to address these challenges, this paper focuses on an experimental evaluation of the impact of DDoS attacks on communications in multiprotocol label switching (MPLS) networks. Additionally, it analyzes the applicability of existing traffic models to accurately reflect the real status of a network. The paper argues that existing network simulation tools such as ns-2 have severe limitations in the context of DDoS attacks and that experimental approaches can help provide a realistic view of network behavior.

The experimental study presented in this paper focused on an isolated environment [6, 16] that reproduced a realistic network, including real carrier-grade Cisco routers, networks, computers systems, protocols and software. Six MPLS network topologies were created and analyzed with regard to the impact of DDoS attacks on MPLS network traffic as well as the functioning of a simulated remotely-controlled power grid. The experiments demonstrate that MPLS virtual private networks (VPNs) alone do not provide proper isolation of virtual circuits, enabling DDoS attacks to severely degrade parallel virtual circuits. The experimental results also show that small changes in network topology can significantly enhance resilience without affecting quality of service (QoS) even when using default router configurations.

2. Problem Statement

Modern critical infrastructures such as power grids, oil and gas pipelines, and water supply systems rely on information and communications technologies for their operation. The advantages of using information and communications technologies include reduced costs as well as greater efficiency, flexibility and interoperability. In the past, critical infrastructure assets were largely isolated and used proprietary hardware and protocols, limiting the threats that could affect them. However, the widespread adoption of commercial-of-the-shelf hardware, software and networking products in modern critical infrastructure assets expose them to disruptive cyber threats.

A DDoS attack on an industrial control system typically engages thousands of infected hosts to flood the victim system with a massive number of packets, consuming network resources and severely reducing communications bandwidth. The result is that that victim system effectively loses its ability to control critical infrastructure assets.

A DDoS attack could ultimately have the same effect as the power grid failure that occurred in Rome on January 2, 2004 [3]. The incident occurred when communications between several remote sites were disabled by a broken water pipe that flooded the server room at a telecommunications service provider, short-circuiting critical hardware. This completely blinded power grid operators, who were unable to monitor or control the remote sites. Fortunately, no additional disturbances occurred during the failure, so the grid remained stable. However, a change in the generated-consumed power balance, possibly caused by weather changes, could have impacted the electrical grid, resulting in a large blackout of the city and affecting other critical infrastructures such as transportation and health care.

Clearly, better protective mechanisms are needed to deal with DDoS attacks, especially in the critical infrastructure, where even short-term outages can have serious consequences. However, there is limited understanding of the effects that DDoS attacks have on real networks, and existing network models and simulation tools are not robust enough to help improve this understanding. For example, existing models [9, 14] often incorporate an abstraction layer that does not account for installation-specific aspects. The basic approach taken by such models is to distribute input traffic on output interfaces by employing a variety of algorithms. But in reality, there are many other factors that influence the distribution of packets, especially in the context of DDoS attacks. For instance, a heavily-loaded router and network might behave very differently from their predicted behavior. Extreme conditions can lead to loss of input and output packets, rendering simplistic input/output traffic counting techniques inapplicable.

In general, there is a lack of models and simulation tools that can accurately reproduce the network state in extreme conditions such as DDoS attacks. The experiments described in this paper reveal that DDoS attacks have serious consequences on remotely-operated critical infrastructure assets even when telecommunications service providers use virtual private networks (VPNs) to isolate traffic. Also, simple topology changes can have a significant impact on network resilience. These changes, which can be deployed using existing routing hardware and software, can help render DDoS attacks ineffective.

3. Experimental Study

Our experiments were designed to explore the consequences of information and communications system disruptions on a simulated power grid. In particular, we considered DDoS attacks on telecommunications services that propagate to and impact the power grid. The focus was on evaluating the effectiveness of existing network models and to demonstrate the impact of small network topology changes on the outcome of DDoS attacks.

3.1. Experimental Approach

The study of complex systems such as modern critical infrastructure assets can be conducted by experimenting with real systems, software simulators or emulators. Unfortunately, for reasons of cost and reproducibility of results, it is difficult to experiment with real systems. Furthermore, if a study seeks to examine the resilience or security of a real system, there are obvious concerns about the potential side effects (faults and disruptions) to mission-critical services.

Software-based simulations can be used very efficiently to study physical systems, primarily because they support low cost, fast and accurate analyses in controlled environments. However, they have limited applicability in the context of cyber security due to the complexity and diversity of real-world networks. Moreover, even when software simulations are able to model network environments, they fail to model network failures with reasonable accuracy [4].

The study described in this paper employed an emulation framework developed in our previous work [6, 16], a modern scientific instrument that helps provide accurate assessments of the impact of cyber attacks on cyber-physical systems used in the critical infrastructure. The emulation testbed, which is based on Emulab [18, 16], provides fidelity, repeatability, measurement accuracy and safety for the cyber layer. The approach is well-established in the field of cyber security [12], and it helps overcome the challenges that arise while attempting to simulate the behavior of information and communications technology components in the presence of attacks or failures.

The emulation testbed engages software simulators to capture the physical components. As mentioned above, simulation is an efficient, safe and low-cost approach that provides rapid and accurate analysis capabilities. While the fidelity is somewhat reduced, software simulation enables disruptive experiments on multiple heterogeneous physical processes. An advantage is that complex models of many important physical systems are described in the literature. These complex models can be simulated to accurately reproduce the behavior of real physical systems. A good example is the power grid, where simulation has become so accurate and trusted that is commonly used to support decision making by electricity transmission system operators.

3.2. Experimental Setup

We recreated the typical architecture of an installation in which a power grid is controlled remotely (Figure 1). Site A runs a simplified model of an energy management system (EMS) [17] to ensure voltage stability. The energy management system continuously monitors and adjusts the operational parameters of the power grid model that runs Site B.

IEEE electrical grid models are extensively used in scientific simulations because they accurately capture the basic characteristics of real power grid infrastructures. We employed the IEEE 39-bus New England system, which incorporates 39 substations and ten generators. The load imposed on the system was based on real data [11]. Intervention by the energy management system was required to keep the grid stable.

To model a real-world environment, it was decided to employ an MPLS network as the telecommunications infrastructure between the energy man-

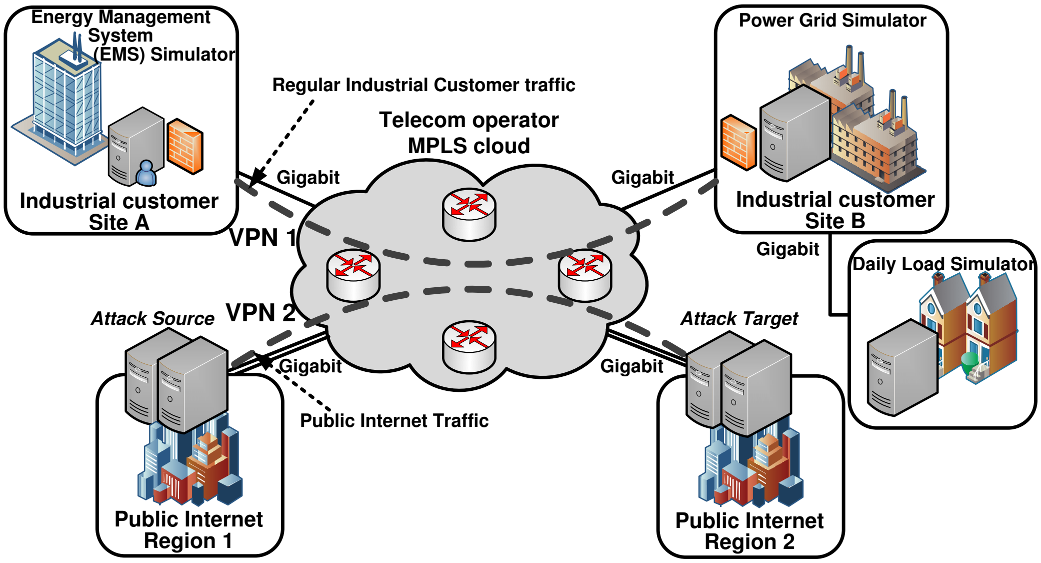

Figure 1: Experimental setup.

agement system and the power grid simulator. MPLS is a protocol that many major telecommunications service providers have used to replace older implementations based on frame relay and ATM. The advantages of MPLS over generic IP networks include more efficient routing capabilities, built-in support for virtual private networks (VPNs) and traffic engineering.

Our Emulab installation was configured to create an MPLS network with four Cisco 6503 routers (figure 1). Two MPLS VPNs were provisioned in this network. VPN 1 acts as a protected virtual circuit between Site A and Site B, an approach that is often adopted by telecommunications service providers to isolate customer traffic. Since a service provider typically routes diverse traffic (including public Internet traffic) through the same MPLS cloud, we use VPN 2 to serve as a virtual circuit between two different "public" regions. The Border Gateway Protocol (BGP) and the Open Shortest Path First (OSPF) Algorithm were used in the MPLS network for traffic routing. Table 1 lists the key equipment used in the MPLS network.

3.3. Telecommunications Disruption

A malicious user can easily manipulate "public" traffic in VPN 2 to induce a telecommunications disruption. To explore the impact of such incidents, we launched a DDoS attack that consumed bandwidth in VPN 2 and measured

Table 1: Key equipment.

Type	**Model**	**No.**	**Description**
Routers	Cisco 6503	4	Cisco carrier-grade routers with four Gigabit experimental interfaces and one control interface; used as MPLS routers.
	Cisco 2911	2	Cisco routers with six Gigabit experimental interfaces and one control interface; used as “public” traffic routers.
Computers	Fujitsu-Siemens	14	Generic personal computers with 2.53 GHz CPUs, 2 GB RAM, two Gigabit experimental interfaces and one control interface; used as experimental nodes.
Switches	Cisco 3750G	3	Cisco switches with 48 ports each; used as a communications network.
Operating System	FreeBSD8.2	–	Used in all the experimental nodes.

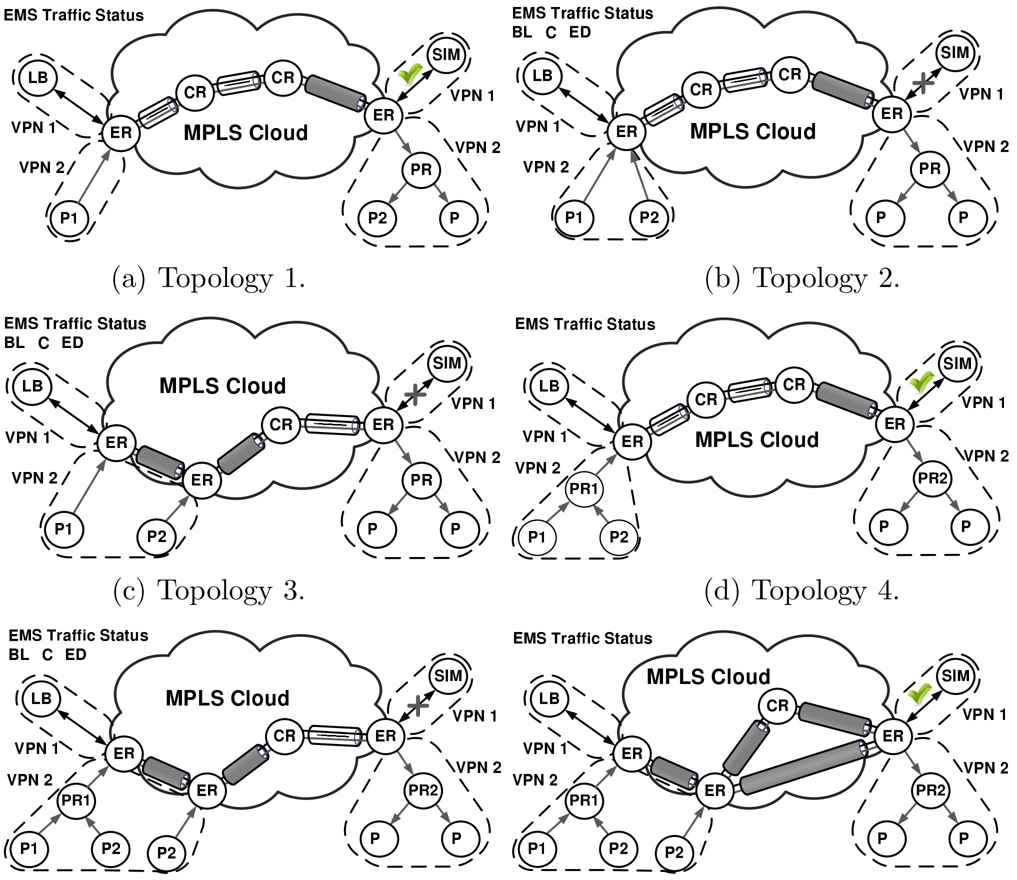

(a) Topology 1.

(b) Topology 2.

(c) Topology 3.

(d) Topology 4.

(e) Topology 5.

(f) Topology 6.

Figure 2: Network topologies used in the experiments.

its effect on the power grid operator's virtual circuit in VPN 1. The attack, which originated from Region 1 and flooded VPN 2 with ICMP packets, was implemented using tools such as TCPReplay and Scapy. The attack throughput at the sender's side (Region 1) for a single node varied between 800 Mbps and 950 Mbps.

3.4. MPLS Network Topologies

Six network topologies were considered in this study (Figure 2). Note that ER denotes an edge router, CR denotes a core router, LB denotes a load balancer, SIM denotes power grid and daily load simulators, PR denotes a public router, and P denotes a public VPN node.

- **Topology 1:** This topology, which is shown in Figure 2(a), incorporates two MPLS core routers and two edge routers. The attack has only one point of origin in VPN 2 and the traffic is merged with the customer traffic in VPN 1, and transmitted over the same network links in the MPLS network.

- **Topology 2:** This topology, which is shown in Figure 2(b), incorporates the same MPLS network as Topology 1. However, a second source of DDoS traffic is added in VPN 2. This leads to a more disruptive effect on the traffic in VPN 1.

- **Topology 3:** This topology, which is shown in Figure 2(c), incorporates one MPLS core router and three edge routers. The second DDoS source in Topology 2 is moved to another edge router in order to test the link dependency of the attack. This is because the outcome of the attack in Topology 2 could be affected by the heavily-loaded edge router.

- **Topology 4:** This topology, which is shown in Figure 2(d), incorporates the same MPLS network as Topology 1 and Topology 2, but the public traffic from two DDoS sources is aggregated in one router (PR1). This topology helps evaluate the effect of external aggregation devices on traffic in the MPLS network.

- **Topology 5:** This topology, which is shown in Figure 2(e), tests the effectiveness of DDoS attacks with aggregated traffic sources. The topology extends Topology 4 using an additional MPLS edge router and an additional DDoS source.

- **Topology 6:** This topology, which is shown in Figure 2(f), tests the effect of multiple connections between MPLS routers on the outcome of DDoS attacks.

4. Experimental Results

This section describes the experimental results. It begins by describing normal operations and proceeds to discuss the effects of DDoS attacks in the six network topologies.

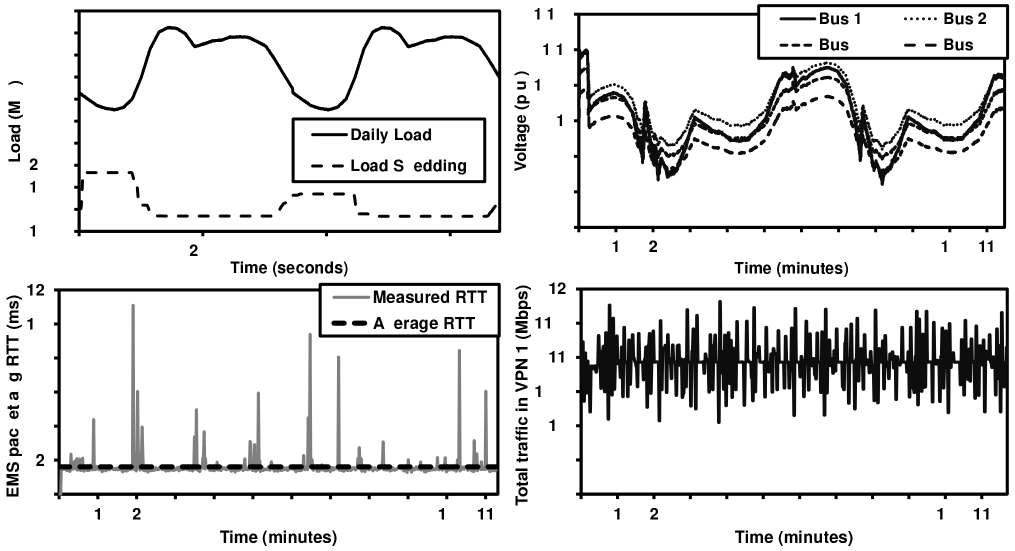

Figure 3: Normal operating state without attack: (a) Daily load and EMS load shedding commands; (b) Voltage status on a selection of five buses; (c) EMS packet round trip time (RTT); and (d) Traffic throughput in VPN 1 including EMS traffic and 10 Mbps UDP-based background traffic.

4.1. Normal Operations

Figures 3(a) and 3(b) show the system in its normal operating state (without an attack). The energy management system is able to keep the grid in a stable state. The voltage levels have small fluctuations due to the daily load values and commands sent by the load shedding algorithm. The energy management system is able to ensure voltage stability and maintain the voltage levels between 0.9 p.u. and 1.1 p.u., which is the range of acceptable operations.

Figure 3(c) shows the network measurements in VPN 1, where the average EMS packet round trip time (RTT) is below 2 ms. This means that the implementation exhibits the operational behavior of a real telecommunications system, where the delivery of high-speed messages must be below the maximum limit of 10 ms according to the IEEE 1646-2004 Standard [8] dealing with communications delays in substation automation implementations. Figure 3(d) shows the performance that is achieved in the presence of 10 Mbps UDP-based background traffic generated using `iperf` [13].

The next six sections describe the effects of DDoS attacks on the telecommunications system and ultimately on voltage stability for the six topologies.

The results are presented in Figure 4, which shows the changes in voltage levels for a selection of four substations.

4.2. Topology 1

The single attack source in this topology (Figure 2(a)) has a limited effect on the traffic in VPN 1 and ultimately no effect on power grid operation (Figure 4(a)). As shown in Figures 5(a) and 5(b) the energy management system is able to communicate with the power grid simulator and the average EMS packet RTT remains below 2 ms. Consequently, voltages are not affected by the attack and the power grid remains in a stable state.

The state of the network can be easily explained by existing traffic models because the cumulative input traffic does not exceed the link capacity. More specifically, the maximum attack throughput of 950 Mbps provides a residual capacity of 50 Mbps, which is enough to handle the 10 Mbps traffic from VPN 1.

4.3. Topology 2

Adding a second attack source to the same edge router (Figure 2(b)) immediately blocks the traffic (Figure 4(b)). The energy management system loses its connection to the power grid simulator, at which point the packet RTT increases to more than one second as shown in Figure 5(a). Shortly after the attack is initiated, the loss of communications with the energy management system leads to instability and voltage collapse (Figure 4(b)). Although the oscillations exhibited by the model cannot be mapped to a real-world scenario, the voltage collapse is a clear indication of grid instability – this could require the power grid operators to rebuild the entire grid. For the purposes of our security study, however, this verifies that the DDoS attack can take the system outside its normal operating limits.

Closer examination of the network measurements in Figures 5(a) and 5(b) reveals that UDP traffic is still able to pass through the network (Figure 5(b)), while the TCP-based EMS traffic stops immediately because no RTTs are reported. This observation could stimulate research focused on DDoS-resilient protocols employing UDP in the critical infrastructure [7].

Existing traffic modeling techniques can give a straightforward answer as to why the TCP traffic was blocked. The reason is that the input flow from the two attack nodes (totaling 1,600 Mbps) exceeds the outgoing link capacity. However, as shown in Figure 5(b), the background traffic is able to pass through despite the presence of two massive DDoS input traffic flows.

An intuitive explanation is that packet losses occur within the router and along the outgoing network link, and also due to network over-provisioning. Since existing network modeling techniques do not incorporate such aspects, this is the first clear example of the importance of the approach adopted in this research.

4.4. Topology 3

The results obtained for this topology confirm the results obtained for Topology 3. The results also highlight the fact that the injection point of the DDoS attack (Figure 2(c)) is not important as long as it intersects the victim's traffic path.

4.5. Topology 4

By aggregating DDoS traffic in one external router (Figure 2(d)), the impact on VPN 1 traffic is reduced and normal remote control operations of the power grid are maintained. However, Figure 5(a) shows that the packet RTT in the energy management system increases from less than 2 ms to an average of 20 ms. This means that get/set commands have a higher latency, although, as shown in Figure 4(d), this does not affect the overall state of the process and the voltages remain stable.

From the network flow point of view, challenges are introduced when applying existing models in this setting. The simple aggregation of attack traffic in the MPLS network causes the cumulative traffic to be equal to the network capacity. Consequently, there should be no more bandwidth left for VPN 1 traffic. But this statement is contradicted by our results, which clearly show that VPN 1 traffic is able to pass through and that bidirectional communications that keep the power grid in a stable state continue to be maintained.

4.6. Topology 5

The results for this topology (Figure 2(e)) confirm our previous analysis. Adding a second DDoS attack source at a different point in the network – even with the aggregation from Topology 4 – completely blocks the energy management system traffic, while a portion of the UDP background traffic is able to pass through (Figures 5(a) and 5(b)). As shown in Figure 4(e), this leads to the loss of communications and ultimately to grid instability.

It is also important to observe that existing network modeling techniques also cannot be applied in this setting. This is because, even with the aggregation of three DDoS traffic sources, UDP traffic still flows in VPN 1.

4.7. Topology 6

This topology extends Topology 5 by incorporating an additional link between the two MPLS edge routers (Figure 2(f)). This allows the ingress edge router to run OSPF load balancing and render the DDoS attack ineffective (Figure 4(f)).

As in the case of Topology 5, using only aggregation-based traffic flow analysis fails to consider network aspects that are crucial in cyber security incidents. Indeed, the reason for introducing Topology 6 was to reinforce the previous results and to clearly show that new network traffic models are needed to model the experimental results. Furthermore, the results for Topology 6 reveal that a network can be made resilient simply by introducing certain architectural changes. These changes can contribute to network deployments that maintain their resilience even when they use default configurations that do not address QoS issues. Nevertheless, it is important to recognize that architectural resilience should be coupled with system configurations that protect against a wider range of attacks.

4.8. Reproducing Results in a Software Simulator

We used the ns-2 software simulator to further defend our statements and to evaluate the impact of DDoS attacks in a completely-simulated environment. The ns-2 discrete event simulator is one of the most popular network simulation tools. Although the ns-3 tool will eventually make ns-2 obsolete, ns-3 is still under development and currently lacks many features that are well-established in ns-2.

First, we evaluated the ability of ns-2 to accurately reproduce the interactions between various UDP-based DDoS traffic sources. We used Topology 5 (Figure 2(e)) for this purpose and generated 150 different traffic settings, where each public node injected traffic ranging from 10 Mbps to 850 Mbps into the MPLS cloud. This scenario underlying Topology 5 was run on the experimental testbed as well as on the ns-2 software simulator.

Topology 5 incorporates three bottlenecks – the external router and the two edge routers – where different traffic sources compete with each other and the routers must ensure the fair selection of packets (because the routers are configured with the default FIFO queuing mechanism). Our objective was to measure the fairness of packet selection from different input queues in the experimental system and to compare them with the results obtained using ns-2. As shown in Figure 6(a), the ns-2 results exhibit large errors compared with the experimental results for default network configurations. Specifically, the errors increase up to 90%, which means that ns-2 has difficulty with traffic fairness. Nevertheless, by tuning the ns-2 parameters (e.g., random packet inter-arrival time) to better reflect reality, the error drops all the way down to 30%. This means that ns-2 can provide better results when its parameters are tuned, a property that has been confirmed by other researchers (see, e.g., [2]).

Next, we evaluated the ability of ns-2 to accurately recreate TCP throughput (i.e., energy management system traffic) in the presence of massive DDoS attacks. For this purpose, we used the previous setting and generated TCP traffic in ns-2 with a similar throughput to the energy management system traffic in the experimental system. In parallel, we generated a DDoS attack and incrementally increased its throughput.

Figure 6(b) demonstrates that, up to a certain threshold (i.e., below an attack rate of 1.52 Gbps), ns-2 exhibits similar behavior as the experimental system. However, in reality, a 10 Mbps increase in DDoD traffice (from 1.51 Gpbs to 1.52 Gbps) can have a dramatic impact on the normal operation of TCP. More specifically, in the experimental system, even a small increase of 10 Mbps reduces TCP throughput from 0.23 Mpbs to 0.04 Mbps that, as shown in the previous sections, can lead to the loss of monitoring and control capabilities, and eventually to power grid instability. Unfortunately, this behavior is not accurately reproduced by ns-2 because TCP throughput remains above 0.12 Mbps for all TCP implementations. This means that, when relying purely on ns-2 and its default configuration, an attack could be interpreted as being ineffective when in fact the opposite is true.

Based on these results, we can conclude that in the context of disruptive cyber attacks (e.g., DDoS attacks), researchers should carefully adjust the parameters of network simulators such as ns-2 to obtain accurate results. Otherwise, as stated by Chertov, *et al.* [4], the results might be open to interpretation and attacks that appear to be ineffective during simulation could have dramatic consequences in reality.

5. Conclusions

The experimental study presented in this paper provides a detailed evaluation of the disruptive effects that DDoS attacks have on information and communications systems and ultimately on cyber-physical systems in the critical infrastructure. In particular, the experimental results reveal that even modern MPLS VPNs with powerful carrier-grade routers can easily fall victim to DDoS attacks.

The results also confirm research by Chertov, *et al.* [4] that existing network simulation techniques and tools such as ns-2 are not well-suited to analyzing computer networks in extreme conditions such as DDoS attacks; there is a great need for new techniques that can incorporate experimental data to approximate network behavior with high fidelity. Finally, the experimental results demonstrate that by engineering architecturally-resilient networks – for example, by employing traffic aggregation together with load balancing – critical infrastructure assets can become more resilient to DDoS attacks even with default security configurations that do not implement QoS.

Our future research will focus on detailed analyses of the behavior of real networks under heavy loads. The ultimate goal is to develop a powerful and robust framework for predicting the behavior of telecommunications networks and other critical infrastructure networks under disruptive DDoS attacks.

IMPORTANT NOTE TO IJCIP TYPESETTERS: I have edited the references in the paper myself. You are welcome to change the order of the references and the order of the items within a reference. However, DO NOT change any capitalization or fonts (e.g., italics) in the references below. Professor Sujeet Shenoi, Editor-in-Chief, IJCIP

References

[1] S. Baker, N. Filipak and K. Timlin, In the Dark: Crucial Industries Confront Cyberattacks, McAfee Report 21900rpt_cip_0311, McAfee, Santa Clara, California, 2011.

[2] M. Bateman, S. Bhatti, G. Bigwood, D. Rehunathan, C. Allison, T. Henderson and D. Miras, A comparison of TCP behavior at high speeds using ns-2 and Linux, *Proceedings of the Eleventh Communications and Networking Simulation Symposium*, pp. 30–37, 2008.

[3] A. Bobbio, G. Bonanni, E. Ciancamerla, R. Clemente, A. Iacomini, M. Minichino, A. Scarlatti, R. Terruggia and E. Zendri, Unavailability of critical SCADA communication links interconnecting a power grid and a telco network, *Reliability Engineering and System Safety*, vol. 95(12), pp. 1345–1357, 2010.

[4] R. Chertov, S. Fahmy and N. Shroff, Fidelity of network simulation and emulation: A case study of TCP-targeted denial-of-service attacks, *ACM Transactions on Modeling and Computer Simulation*, vol. 19(1), pp. 4:1–4:29, 2008.

[5] C. Deccio, Maintenance, mishaps and mending in deployments of the domain name system security extensions (DNSSEC), *International Journal of Critical Infrastructure Protection*, vol. 5(2), pp. 98–103, 2012.

[6] B. Genge, C. Siaterlis, I. Nai Fovino and M. Masera, A cyber-physical experimentation environment for the security analysis of networked industrial control systems, *Computers and Electrical Engineering*, vol. 38(5), pp. 1146–1161, 2012.

[7] P. Haller and L. Marton, Prediction and congestion control algorithm for networked motion tracking, *Control Engineering Practice*, vol. 17(11), pp. 1265–1272, 2009.

[8] Institute of Electrical and Electronics Engineers, IEEE 1646-2004 Standard: Communication Delivery Time Performance Requirements for Electric Power Substation Automation, New York, 2004.

[9] A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E. Kolaczyk and N. Taft, Structural analysis of network traffic flows, *Proceedings of the International Joint Conference on Measurement and Modeling of Computer Systems*, pp. 61–72, 2004.

[10] J. Leyden, Biggest DDoS attack in history hammers Spamhaus, *The Register*, March 27, 2013.

[11] M. Manera and A. Marzullo, Modeling the load curve of aggregate electricity consumption using principal components, *Environmental Modeling and Software*, vol. 20(11), pp. 1389–1400, 2005.

[12] J. Mirkovic, T. Benzel, T. Faber, R. Braden, J. Wroclawski and S. Schwab, The DETER Project: Advancing the science of cyber security experimentation and test, *Proceedings of the IEEE International Conference on Technologies for Homeland Security*, 2010.

[13] National Center for Supercomputing Applications, Iperf: The TCP/UDP Bandwidth Measurement Tool, University of Illinois at Urbana-Champaign, Urbana, Illinois (`iperf.fr`), 2011.

[14] A. Pascale and M. Nicoli, Adaptive Bayesian network for traffic flow prediction, *Proceedings of the IEEE Statistical Signal Processing Workshop*, pp. 177–180, 2011.

[15] S. Rahimi and M. Zargham, Analysis of the security of VPN configurations in industrial control environments, *International Journal of Critical Infrastructure Protection*, vol. 5(1), pp. 3–13, 2012.

[16] C. Siaterlis, A. Garcia and B. Genge, On the use of Emulab testbeds for scientifically rigorous experiments, *IEEE Communications Surveys and Tutorials*, vol. PP(99), pp. 1–14, 2012.

[17] T. Tuan, J. Fandino, N. Hadjsaid, J. Sabonnadiere and H. Vu, Emergency load shedding to avoid risks of voltage instability using indicators, *IEEE Transactions on Power Systems*, vol. 9(1), pp. 341–351, 1994.

[18] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb and A. Joglekar, An integrated experimental environment for distributed systems and networks, *ACM SIGOPS Operating Systems Review*, vol. 36(SI), pp. 255–270, 2002.

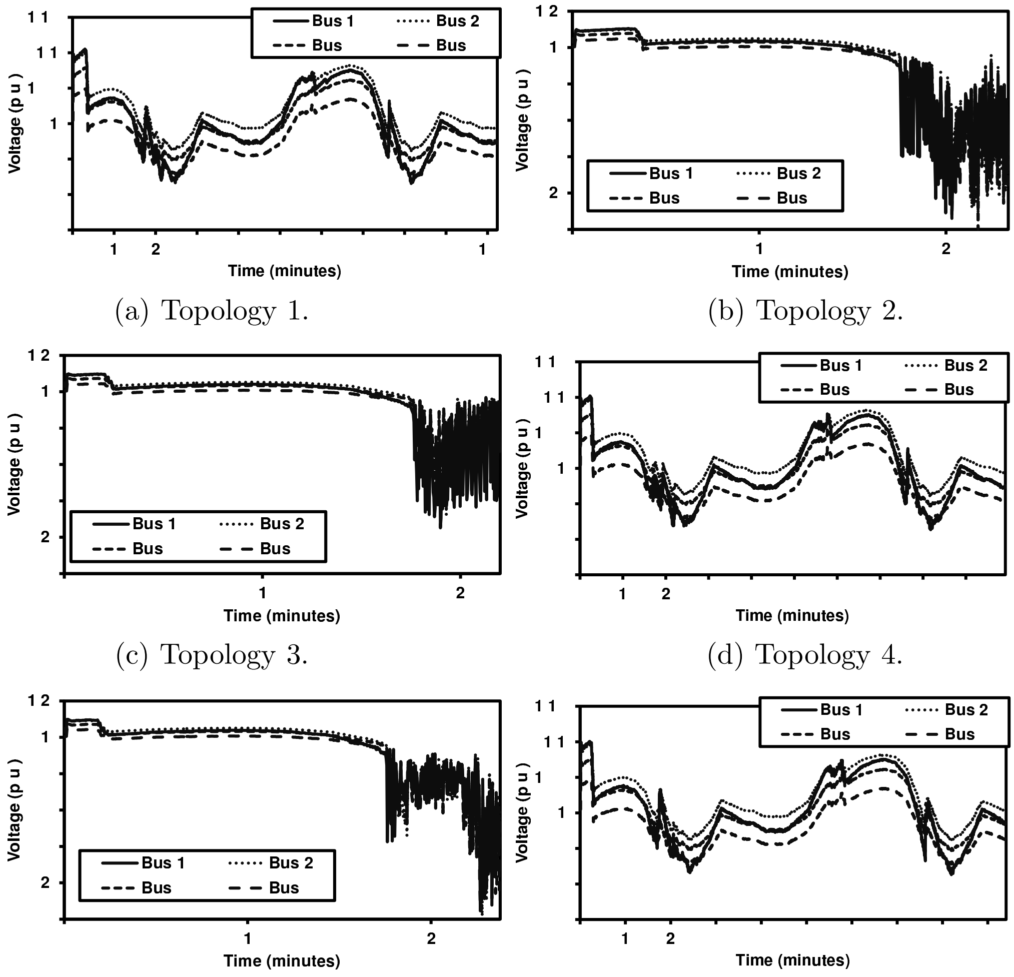

(a) Topology 1.

(b) Topology 2.

(c) Topology 3.

(d) Topology 4.

(e) Topology 5.

(f) Topology 6.

Figure 4: Effects of DDoS attacks on voltage stability.

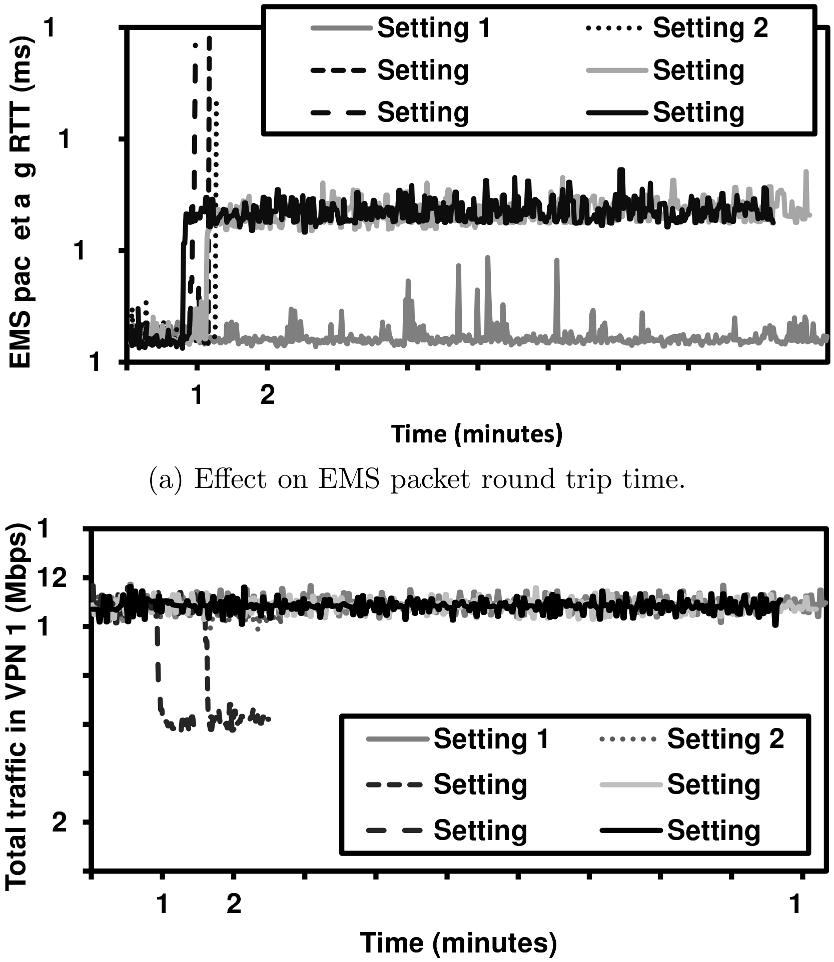

(a) Effect on EMS packet round trip time.

(b) Effect on VPN1 traffic throughput.

Figure 5: Effects of DDoS attacks on VPN1 traffic.

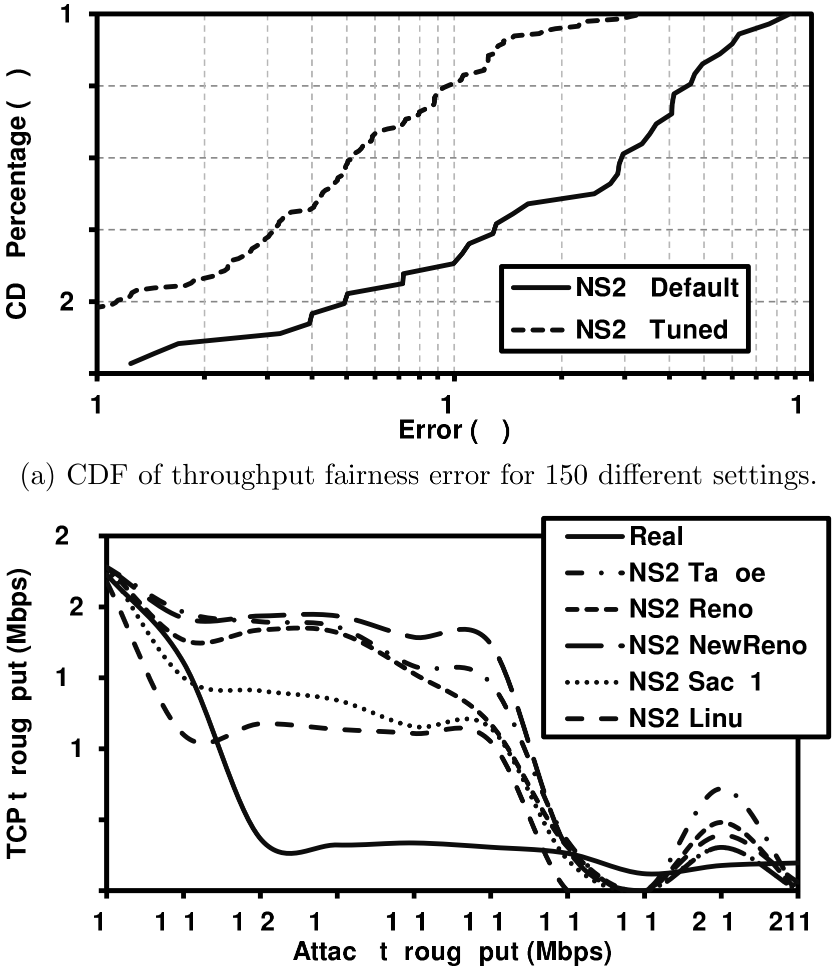

(a) CDF of throughput fairness error for 150 different settings.

(b) TCP (EMS) traffic throughput in the presence of a DDoS attack.

Figure 6: Experimental results obtained with ns-2.