PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Optical mitigation of DDoS attacks using silicon photonic switches

Shen, Yiwen, Goodfellow, Ryan, Strom Glick, Madeline, Bartlett, Genevieve, Bergman, Keren

Yiwen Shen, Ryan Goodfellow, Madeline Strom Glick, Genevieve Bartlett, Keren Bergman, "Optical mitigation of DDoS attacks using silicon photonic switches," Proc. SPIE 11308, Metro and Data Center Optical Networks and Short-Reach Links III, 113080J (31 January 2020); doi: 10.1117/12.2549080

SPIE.

Event: SPIE OPTO, 2020, San Francisco, California, United States

Invited Paper

Optical Mitigation of DDoS Attacks using Silicon Photonic Switches

Yiwen Shen[a], Ryan Goodfellow[b], Madeleine Strom Glick[a], Genevieve Bartlett[b], and Keren Bergman[a]

[a]Columbia University, New York, NY, USA
[b]Information Sciences Institute, Marina del Rey, CA, USA

ABSTRACT

In this paper, we demonstrate the integration of a SiP switching platform to improve real-world Distributed Denial of Service (DDoS) defense systems. We demonstrate how DDoS mitigation in the optical domain can be transparent to network and application layers, allowing for reconfiguration and tuning. Additionally, we show how optical domain DoS mitigation provides significant cost reduction—with a 1/3 cost reduction—compared to traditional mitigation using electronic counterparts. Our approach is ideal for data-center deployments, and our testbed topology mirrors a standard data center set up.

Keywords: Cybersecurity, Data Center, DDoS, Photonic Switching, Micro Ring Resonator

1. INTRODUCTION

A volumetric denial of service attack (DoS) overwhelms a victim service by sending a high volume of network traffic, inundating hardware near and at the victim service. Most commonly, these large-scale attacks are in the form of a distributed attack (DDoS) from thousands to tens of thousands of hosts, often co-opted into a botnet through malware. These attacks cause hundreds of thousands of dollars in lost revenue for small to medium businesses, and millions of dollars of damage for large enterprises. In 2019, the *average* cost to an enterprise in lost revenue, control measures and damage repair per attack was over $2 million in U.S. currency.[1] A broad range of research efforts over the past two decades have worked to address detection, characterization, classification and mitigation of these attacks,[2–5] but DDoS continues to be a serious economic and technical threat.

At line-rates, traditional mitigation over high-bandwidth electronic packet-switched domains is complex and costly, often leaving smaller sites unable to provide adequate protection for their hosted services. While cloud services use replication across multiple sites to limit the effects of DoS on a particular application, attacks still target cloud infrastructures, incurring costs that are ultimately shouldered by cloud customers. As organizations trend towards cloud-hosted Internet services and applications, the relatively small group of hyperscale data centers that comprise the vast majority of the cloud market are taking on a larger and larger attack surface that increases as these data centers host new customers and applications.

In this paper we present a mitigation approach based in the optical domain, which offers a simpler and less costly solution than traditional mitigation with electronic packet-switching. In the optical domain, the incorporation of a silicon photonic switching platform offers a cost-effective solution that can selectively route individual wavelengths of high-bandwidth WDM links, enabling timely and efficient mitigation decisions that are transparent to higher network layers. We demonstrate integration of microring-resonator-based silicon photonic (SiP) switches within a data center testbed to defend against DDoS attacks. This platform targets critically important real-world defense systems by significantly reducing complexity through the use of transparent SiP switches. Our approach is designed for standard data center topologies and is adapted to the low-port counts of SiP switches. We show how this system not only has the potential for cost reductions—by a third—compared to electronic counterparts, but also demonstrate how this mitigation can be made completely transparent to higher network layer protocols and client applications by providing network operators with access to network topology reconfiguration of the optical physical layer.

Further author information: (Send correspondence to R.G.)
R.G.: E-mail: rgoodfel@isi.edu

Metro and Data Center Optical Networks and Short-Reach Links III, edited by Atul K. Srivastava, Madeleine Glick, Youichi Akasaka, Proc. of SPIE Vol. 11308, 113080J · © 2020 SPIE
CCC code: 0277-786X/20/$21 · doi: 10.1117/12.2549080

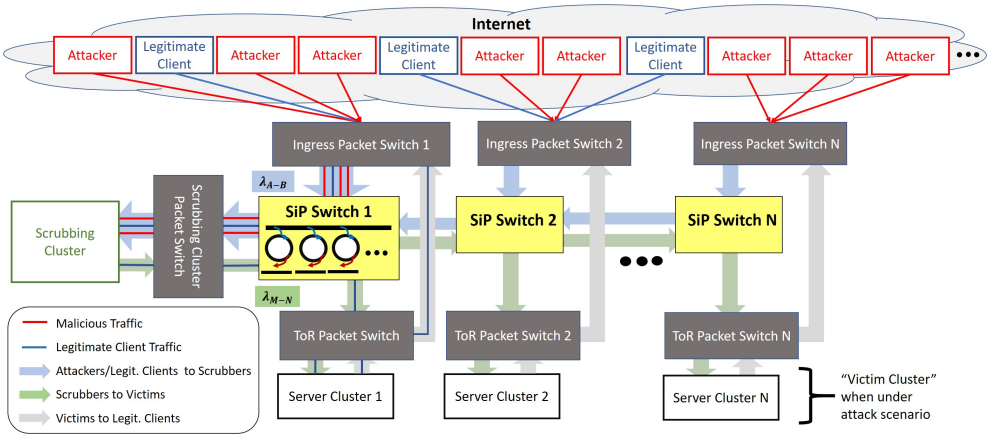

Figure 1. System architecture (under an attack scenario)

2. SYSTEM ARCHITECTURE

Figure 1 illustrates our system architecture. External malicious and legitimate client nodes communicate with the internal data center (DC) network that consists of the ingress electronic packet switches (EPSs), the scrubbing EPS and victim cluster ToR EPS. This is a common real-world setup where quick detection leads to a majority of the attack being black-holed (discarded). The leftover traffic remains a mix of attack and non-attack and must be processed by the scrubbing cluster which removes the remaining malicious packets. A detection system notifies the control plane which ingress EPS ports are likely to be receiving malicious traffic, upon which the platform presented in this paper reacts. The silicon photonic (SiP) switches are used to facilitate reconfiguration of the physical network topology between these clusters depending on whether or not the system is defending against a DDoS attack.

During an attack, the ingress traffic over any number of uplinks is a mix of legitimate client traffic and attack traffic destined to the DDoS victim. After the detection system identifies the input ports likely to contain the attack, the control plane directs traffic from the identified ports to specific output ports transmitting known wavelengths (λ_{A-B}) to the silicon photonic switch. Each individual wavelength carries a mixture of malicious and legitimate clients' traffic channels. The silicon photonic (SiP) switches consist of microring resonators (MRRs) that receive desired signals by tuning their resonance response to the specific wavelength of the signal. When the signals from the ingress EPS carrying malicious and legitimate clients' traffic reaches the SiP switch, the MRRs whose output links are physically connected to the scrubbing cluster's EPS tune to these wavelengths, i.e. λ_{A-B}, redirecting them to the scrubbing cluster instead of the victim cluster, transparent to the intended destination contained in the header. The scrubbing cluster filters the received traffic and removes packets from malicious source nodes and forwards the remaining legitimate client's traffic towards the SiP switch on wavelengths λ_{M-N}. These are directed by the SiP MRRs to the victim cluster ToR EPS, allowing the victim cluster to receive legitimate clients' traffic. From there, the victim cluster uses a static link that joins the victim cluster ToR EPS to the ingress EPS to reply only to the legitimate clients, and the requests from malicious nodes are ignored.

Under normal operation when there is no attack, the MRRs of the SiP switch directs traffic output from λ_{A-B} from clients directly to the so-called victim cluster ToR EPS, thereby bypassing the scrubbing cluster, and the victim cluster communicates back to the clients using the static link between the victim cluster and ingress EPS.

2.1 Scalability

The scalability of this platform is defined through the number of MRRs required for handling each wavelength and redirecting it to the scrubber. The scrubber nodes are a limited resource as they are employed for removing leftover malicious traffic after the majority of traffic has already been black-holed. Each individual scrubber node handles a single output wavelength (containing a mixture of multiple attack and non-attack traffic channels) from

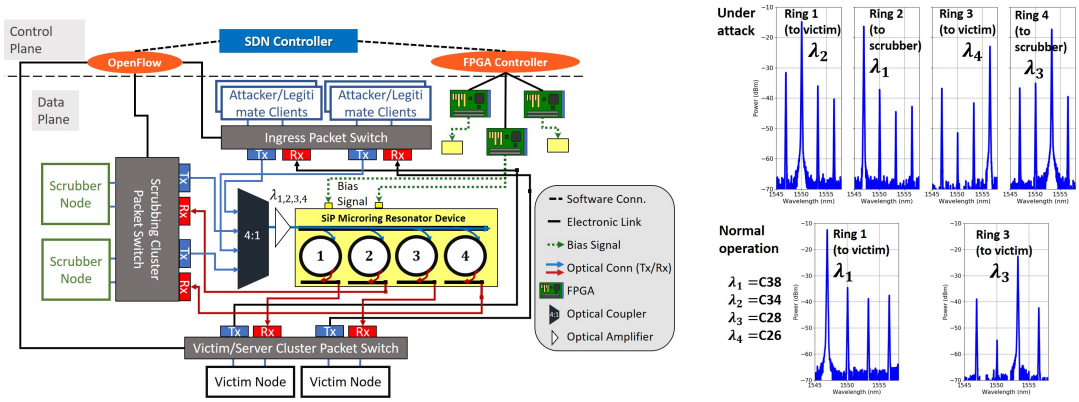

(a) (b)

Figure 2. (a) Network architecture showing connections between the control plane and data plane, (b) output spectrum of each MRR showing the resonance wavelength during normal operation and when under attack

the ingress EPS. Under an attack scenario, two MRRs are required per output wavelength - one MRR for directing traffic from the ingress EPS to the scrubber EPS, and another to direct traffic from the scrubber EPS to the victim cluster ToR EPS. When not under attack, only a single MRR is required per wavelength for directing the traffic from the legitimate clients to the so-called victim. Therefore, although SiP switches have typically low port counts, the number of MRRs required will not limit system scalability. The SiP switches are easily integratable above the Edge/ToR packet switches to protect the servers in each rack within the network topology, as depicted in Fig. 1.

3. TESTBED SETUP

We built a testbed to test the feasibility of our concept (Figure 2(a)). An SDN controller is used to manage all active components of the data plane, which consists of manipulating the flow tables on the EPSs through OpenFlow, and controlling the state of the SiP switch using an FPGA interface (the details are described in [6]). At the data plane layer, a cluster of 4 servers are used to represent the attacker and legitimate client nodes, while another two servers are used to represent the scrubber and victim nodes each, respectively. The ingress, scrubbing, and victim cluster EPSs are virtually partitioned from two PICA8 Ethernet switches. Fig. 2(b) shows the optical spectrum seen by the receiving side of the commercial 10G SFP+ transceivers after traveling through one of four MRRs. The highest peak is the wavelength that the ring is tuned to, while other peaks are crosstalk with at least 15 dB between the two highest peaks.

3.1 Experiment Workflow

We carried out a DDoS attack using multiple attack nodes and flooded a victim node with *hping* (an open-source DDoS attack tool) while legitimate client nodes sent traffic to the victim nodes using *Iperf* (an open-source traffic generator). The SDN controller adds Layer-2 flow rules to direct all the traffic from the attack and legitimate clients to the output port with the transceiver transmitting λ_1 and λ_3. Ring 2 and 4 of the SiP switch whose output links are connected to the scrubbing cluster EPS are tuned to λ_1 and λ_3 respectively, so that the mix of attack and legitimate traffic reaches the scrubber nodes. Each scrubber has two 10G network interfaces, and a virtual network bridge connects them to allow Layer-2 forwarding of the input traffic from one interface to the other. In between the forwarding, the scrubber will black-hole the attack traffic (from *hping*), which was done by first having the attack nodes send packets with an arbitrarily chosen TTL value, and using the *iptables* utility program to set a rule that drops all packets with this TTL value. The rest of the traffic (*Iperf* from the legitimate clients) was then forwarded back out on λ_2 and λ_4. Ring 1 and 3 of the SiP switch whose output link connects to the victim cluster are tuned to λ_2 and λ_4, respectively, so that the victim nodes receive the legitimate clients' traffic, and are able to reply to the legitimate clients through a static electronic connection to the ingress EPS.

Electronic Packet Switch Model

$$n_{Eswitch} = \frac{n_{pod} + n_{ingress}}{radix}$$
$$n_{transceiver} = 2 \times (n_{ingress} + n_{pod} + n_{switch})$$
$$n_{cable} = n_{pod} + n_{ingress} + n_{switch}$$

Silicon Photonic Switch Model

$$n_{Oswitch} = \frac{n_{pod}}{n_{\lambda}}$$
$$n_{transceiver} = n_{switch} + n_{pod}$$
$$n_{cable} = n_{pod} + n_{ingress} + n_{switch}$$

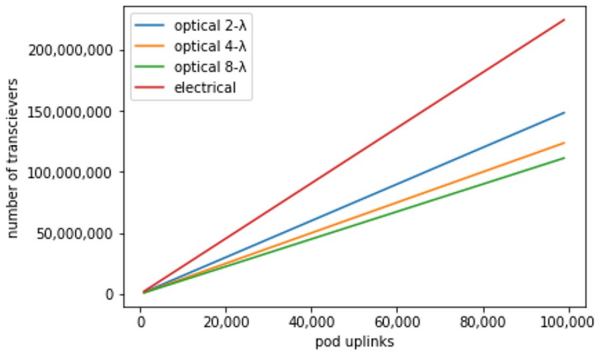

Figure 3. Estimated cost model for 100k uplinks assuming transceivers: $300 and fiber trunks: $50

Under normal operation (not under attack), the SiP switch is configured with Ring 1 tuned to λ_1 and Ring 3 tuned to λ_3, so that the traffic directly flows from the legitimate clients to the victim nodes and back through the static links.

4. BENEFITS TO COST AND COMPLEXITY

The provisioning of dynamic network paths to reroute suspect traffic to the scrubbing cluster is accomplished through an intermediary routing device between the ingress and ToR switches. We compare the cost of the using a conventional EPS and conventional electro-optical network links with our SiP platform on a varying number of wavelengths, for a data center with up to 100k client cluster (also referred to as pods) uplinks (Fig. 3). It is clear that the cost for the SiP model grows much slower than for the electronic case, with the model showing that at two wavelengths, the cost of the photonic strategy scales approximately 1/3 slower rate than the electronic counterpart. This is due to the fundamental capability of the SiP platform to selectively route individual wavelengths within a WDM link without the need for OEO conversion. The primary cost saving in replacing the EPS with a SiP switch is in the cost of intermediate transceivers and cabling.

Potentially more importantly, traditional packet-switched defense strategies involve covertly rerouting traffic through a mitigation network which, requires complex mechanisms at Layer-2 or Layer-3 that in many ways go against how networks at these layers are designed to work. A controllable optical platform provides this capability at Layer-1, where switching is performed through physical topological reconfiguration transparent to higher network layers (and subsequently to applications), thus removing the complexity required for such covert actions. This is especially beneficial for large complex networks where subtle, hard-to-track miscues can have catastrophic effects.

Current solutions for defending against DDoS attacks employing scrubbing clusters use packet switches and require transceivers and cabling for each individual link, which is expensive and requires high wiring complexity. Replacing this hardware with the SiP switch, which can *combine connections using wavelength-division multiplexing (WDM)* and redirect desired traffic channels using high-bandwidth wavelength routing capabilities, in addition to lower cost and small footprint, is a vast improvement.

From a cost perspective the main trade off is exchanging an electronic packet switch and the optical transceivers required for optical to electrical and then electrical to optical conversion for a transparent optical switch. For the cost of the silicon photonic switch chip, we estimate a cost of less than 700 dollars for a switch with four to eight channels. This cost which includes electronic drivers and packaging is based on current fabrication and packaging costs and assumes a 0.1 factor drop in price from prototype to moderate volume production. With four inputs and outputs as in Fig. 2, we would require eight optical transceivers.

Consider Figure 4 which compares $\emptyset$) no mitigation, e) electronic mitigation, λ_*) silicon photonic mitigation strategies for a standard pod-based data center architecture. The model shows that at $\lambda = 2$ the cost of the photonic strategy scales at close to a 40% slower rate than the electronic strategy.

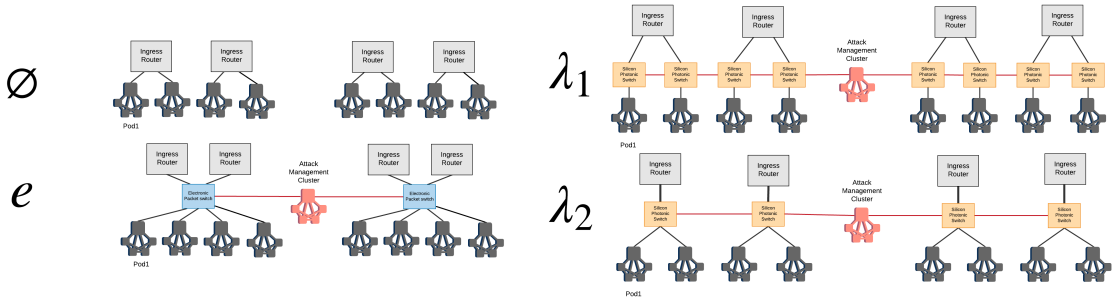

Figure 4. Alternative SiP Switch Integrations

From a system level perspective, Data center traffic typically comes in through an ingress router, goes through a firewall to the destination pod hosting the virtual machine the traffic is destined for. When dedicated attack management resources are present, dynamic network paths must be created to reroute suspect traffic to mitigation resources. To accomplish this, an intermediary routing device must be placed between the ingress routers and pod spine switches. Consider the cost of doing so with conventional electro-optical network links in use today and electronic packet switches, this cost is represented in the electronic switch model (e). The primary variable here is the radix of the routing switches. Alternately, in the optical switch model the primary variable of interest is the wavelength capacity of the switch uplinks to the ingress routers. In both cases the larger the radix/lambda-capacity the smaller the number of switches required. In the optical switch case, there are also fewer transceivers required as there is no OEO conversion required.

4.1 Benefits to System Complexity

In this section we compare the complexity trade-offs of transiting suspect traffic through a scrubbing cluster at they physical layer (1), link layer (2) and network layer (3). At each of these layers we have the same problem - transparently routing a select subset of traffic through a scrubbing cluster on its way to a final client destination. The mechanisms at each layer however, are quite different and present significantly different complexity boundaries from a network management, automation and configuration perspective.

Routing traffic through a scrubbing cluster at the physical layer is described in Sections 2 and 3 and is depicted in the diagram on the left side of Figure 5. By taking on the problem at this layer the only setup that is needed is at the ingress router and the SiP switch. The ingress router is configured by an out of band control plane that is informed by the data centers attack detection systems and the SiP switch is configured to route the set of lambdas corresponding to suspect traffic to the scrubbing cluster and everything else goes to the client cluster as usual. The complexity of network configuration and management here is contained to the configuration of the ingress router and the very simple setup of the SiP switch.

Now consider taking on the problem at layer 2 e.g. the ingress router and client machines have IP addresses but everything in between is governed by link layer switching. In this case the mechanism of isolating suspect traffic is virtual local area networks (VLAN)*. The ingress router now adds a VLAN tag to suspect traffic and the electronic packet switches in the core of the network must ensure that the forwarding tables associated with this designated VLAN push traffic through the scrubbing cluster. Furthermore they must maintain an overlapping forwarding table for legitimate traffic that are off the designated VLAN to traffic that is either a) not tagged by the ingress router or, b) passed through the scrubbing cluster as benign and effectively untagged - still transits to the appropriate clients. This extra tag stripping could take place in the scrubbing cluster itself forcing it to be VLAN aware which is a bit of a complexity leak, or through directional filtering rules at some switch between the scrubbing cluster and client clusters. In any case the maintenance of the network is higher due to layer 2 considerations for transparent routing through the scrubbing cluster.

Finally consider taking on the problem at layer 3 e.g. the client clusters and the scrubbing clusters have their own routers which connect to the ingress router, potentially through intermediary routers. In this example

*clearly there are other layer 2 strategies, VLANs are selected as an example

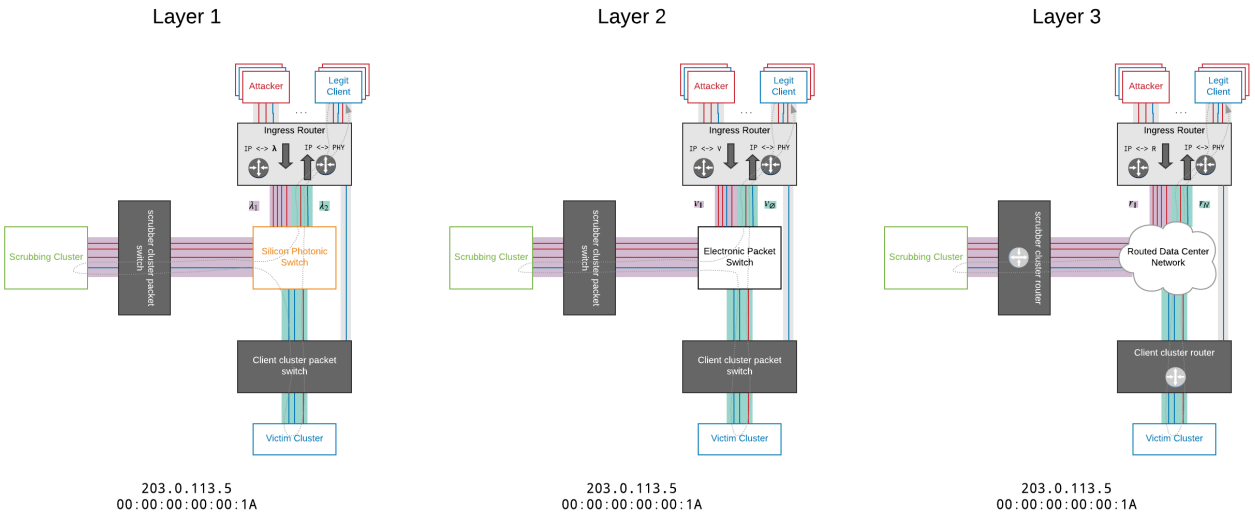

Figure 5. Routing suspect traffic at different layers. The diagram on the left shows routing using SiP switches at layer one. The diagram in the middle shows routing using electronic packet switches using VLANs at layer 2. The diagram on the right shows the use of routers at layer 3. Example layer 3 and 2 addresses are depicted at the bottom of each diagram respectively.

we assume that the network is based on the border gateway protocol (BGP). Thus the mechanism of isolating suspect traffic is through advertisements. In some routed data center isolation systems such as Ethernet Virtual Private Network (EVPN) advertisements are combined with and overlay/underlay encapsulation barrier such as VXLAN to ensure VLAN-like traffic separation. In these systems the routers connected to the client clusters and the scrubbing cluster must create advertisements on behalf of the clients behind them. To achieve the desired effect the ingress router applies a special routing target to suspect incoming traffic. The client and scrubbing cluster routers can export advertisements on behalf of the same client addresses, but with different route targets so that suspect traffic always goes through the scrubbing cluster. Similar to the layer 2 scenario somewhere between the scrubbing cluster and the client cluster, a network component needs to make the translation in routing targets between traffic that ingresses as suspect, is cleared by the cluster and egresses as benign on a different routing target. On EVPN based systems distinct virtual network identifier (VNI) tags could also be used for this purpose but still require a intermediary that is aware of the network scheme in play to translate between tags.

5. CONCLUSION

In this paper we have presented a technique for mitigating DoS attacks on data center networks using silicon photonic switches. We have demonstrated the viability of the approach using a network testbed with real hardware. Our calculations show that mitigation of hostile traffic at layer 1 using SiP switches can bring significant cost savings over traditional electronic packet switch based techniques. We also analyze the system level complexity of hostile traffic mitigation approaches at multiple layers of the network stack and conclude that there are beneficial trade-offs in favor of mitigation at layer one that has a simplifying effect on overall system operation and management.

REFERENCES

[1] Netscout, *NETSCOUT Threat Intelligence Report - 1H 2019* (2019 (accessed Dec 2019)). https://www.netscout.com/threatreport.

[2] Yan, Q., Yu, F. R., Gong, Q., and Li, J., "Software-defined networking (sdn) and distributed denial of service (ddos) attacks in cloud computing environments: A survey, some research issues, and challenges," *IEEE Communications Surveys & Tutorials* **18**(1), 602–622 (2015).

[3] Zargar, S. T., Joshi, J., and Tipper, D., "A survey of defense mechanisms against distributed denial of service (ddos) flooding attacks," *IEEE communications surveys & tutorials* **15**(4), 2046–2069 (2013).

[4] Peng, T., Leckie, C., and Ramamohanarao, K., "Survey of network-based defense mechanisms countering the dos and ddos problems," *ACM Computing Surveys (CSUR)* **39**(1), 3 (2007).

[5] Mirkovic, J. and Reiher, P., "A taxonomy of ddos attack and ddos defense mechanisms," *ACM SIGCOMM Computer Communication Review* **34**(2), 39–53 (2004).

[6] Shen, Y., Gazman, A., Zhu, Z., The, M. Y., Hattink, M., Rumley, S., Samadi, P., and Bergman, K., "Autonomous dynamic bandwidth steering with silicon photonic-based wavelength and spatial switching for datacom networks," in [*2018 Optical Fiber Communications Conference and Exposition (OFC)*], 1–3, IEEE (2018).