

Fighting Fire with Light: Tackling Extreme Terabit DDoS Using Programmable Optics

Matthew Nance Hall[†], Guyue Liu[‡], Ramakrishnan Durairajan[†], Vyas Sekar[‡]

[†] University of Oregon, [‡] Carnegie Mellon University

ABSTRACT

Distributed denial-of-service (DDoS) attacks are a clear and present threat to both today’s and future network infrastructures. Attacks are constantly growing in sophistication with new threats emerging and likely amplified with other technology trends (e.g., amplification, IoT botnets, 5G connectivity). While great progress has been made in devising many types of mitigation strategies, they are found wanting in light of advanced large-scale attacks and our ability to minimize the impact of the attacks on legitimate services.

In this work, we explore a new opportunity for bolstering our DDoS defense arsenal by leveraging recent advances in programmable optics. We envision ONSET: an Optics-enabled In-Network defenSe for Extreme Terabit DDoS attacks. Our approach seeks to isolate and steer attack traffic by dynamic reconfiguration of (backup) wavelengths. This physical isolation of attack traffic enables finer-grained handling of suspicious flows and offers better performance for legitimate traffic in the face of large-scale attacks. In this position paper, we demonstrate the preliminary promise of this vision and identify several open problems at the intersection of security, optical, and systems communities.

CCS CONCEPTS

• Networks → Programmable networks; Physical links; • Security and privacy → Denial-of-service attacks;

KEYWORDS

Terabit DDoS defense, Programmable optics

ACM Reference Format:

Matthew Nance Hall[†], Guyue Liu[‡], Ramakrishnan Durairajan[†], Vyas Sekar[‡]. 2020. Fighting Fire with Light: Tackling Extreme Terabit DDoS Using Programmable Optics. In *Workshop on Secure Programmable Network Infrastructure (SPIN’20)*, August 14, 2020, Virtual Event, NY, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3405669.3405824>

1 INTRODUCTION

Distributed denial-of-service (DDoS) attacks are on the rise [1, 9, 21, 43]. The immense attack volumes, attack diversity, sophisticated attack strategies, and the low cost of launching them make DDoS

attacks a critical cybersecurity issue in today’s and future Internet infrastructure. DDoS defense is not a new problem and prior work has made significant progress in devising mitigation strategies to tackle DDoS attacks. These range from packet scrubbing solutions [5, 7, 13, 16, 35], in-network filtering [8, 20, 29, 52], to more recent routing around congestion methods [47], as well SDN/NFV-based elastic defenses [19, 53]. Despite these advances, the rise of new-age extreme terabit attacks mandates a critical rethinking of DDoS defense strategies.

In this work, we pursue a new opportunity for bolstering our DDoS defense arsenal by leveraging advances in programmable optics. A specific *technology push* we propose to exploit here is the ability to program a subset—not all—of the optical components. For example, ROADMs enable steering of existing wavelengths at finer granularities and enable rerouting traffic on the order of sub-seconds at the physical layer [37]. Another example is the improvements to amplifier modeling [27], pointing towards a rapidly programmable backbone in the near future. These timescales and programmability are recent advances that are being deployed as commodity solutions.

While programmable optics have been used for traditional network management tasks such traffic engineering and routing [17, 22, 42, 46], its benefits have not been explored in the context of DDoS defense. As a starting point, we envision two use cases where optics-enabled defenses can provide significant benefits. First, consider rerouting traffic to “packet scrubbers” to detect and drop malicious packets. Unfortunately, benign traffic will suffer performance impacts and this rerouting can induce additional network congestion. With an optical-enabled defense, such congestion problems can be tackled using *dynamic capacities* (i.e., scaling the capacity of optical paths on-demand) by opportunistic reconfiguration of wavelengths. Second, if we can combine it with some lightweight mechanism to detect suspicious vs. benign traffic in the data plane, legitimate traffic can be routed without congestion on *isolated* wavelengths. This has the dual benefit of reducing the total scrubbing cost and reducing the performance impact on benign traffic.

Motivated by these use cases, we make a case for the ONSET framework—Optics-enabled In-Network defenSe for Extreme Terabit DDoS attacks. Our vision in ONSET is to bridge packet and optical networks to provide new defense capabilities that would otherwise be impractical. In this paper, we identify two key classes of technical challenges in realizing this vision and sketch a preliminary roadmap to address them:

- (i) *Models for ONSET*: While the above use cases are instructive, we need to systematically understand if, when, and how specific optical-layer opportunities can help different DDoS scenarios. To

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SPIN’20, August 14, 2020, Virtual Event, NY, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8041-6/20/08...\$15.00

<https://doi.org/10.1145/3405669.3405824>

this end, we present a simple throughput model for a direct attack with a classical packet scrubber deployment for mitigation and demonstrate the efficacy of ONSET. By isolating the traffic into two groups (*i.e.*, suspicious and trusted), we show that ONSET can reroute the suspicious traffic to a scrubber and send trusted traffic directly to the destination—improving throughput by 25% to 51%, while reducing latency by 33% to 65%.

- (ii) *Design of ONSET*: Motivated by these analytical results, we present the design goals and system requirements of ONSET, building atop ONOS [3] SDN controller. Our design includes a (a) flexible data plane with capabilities and APIs to reconfigure wavelengths in the face of ongoing traffic and to coordinate optical components and programmable switches, (b) robust control plane that isolates/steers suspicious flows into separate wavelengths in a performance-aware manner and runs advanced defenses for extreme terabit attacks, and (c) new suite of applications for network operators.

Looking beyond these specific scenarios, we speculate that the ONSET approach can be more broadly applied to other security issues. For example, advanced attackers can map and launch sustainable, adaptive attacks by fingerprinting the network topologies, services, and even defenses. With novel optical features such as rapid wavelength reconfiguration and amplifier tuning, however, we can introduce new dynamic adaptation capabilities to combat such advanced DDoS attacks. For example, wavelengths can be changed at different locations of the network to thwart ongoing network reconnaissance [39]. Similarly, wavelengths can also be strategically added on-demand to create new capacity to alleviate attack-induced congestion for advanced infrastructure (e.g., Crossfire) attacks [32]. We hope this position paper spurs further conversation and efforts in the programmable networks community.

2 A CASE FOR OPTICS-ENABLED DDOS DEFENSE

DDoS attacks keep escalating in volume, diversity, and dynamism. For example, new types of link flooding attacks [32] target at the transit core of the Internet, which not only impacts a few endpoints but all users of that link. Motivated by this evolving attack landscape, researchers have proposed two new types of “agile” defenses: (1) using NFV-based middleboxes [19] to replace traditional hardware appliances and to launch flexible defenses based on the scale and type of attack; and (2) by leveraging programmable data planes [53] to defend attacks rapidly at the infrastructure level.

While these defenses are as compelling as ever, they treat the underlying topology (*e.g.*, optical layer) as a static entity and, in the worst case, are oblivious to the topology. One might ask “why are not all optical links loaded to maximum capacity at all times?” This question trivializes capacity planning and infrastructure design due to two false assumptions about the optical layer. First, operators can easily provision all fiber links in their network to maximum capacity at a negligible cost. While it is true that a strand of fiber has a fundamental limit to the amount of data it can carry (known as the Shannon limit [18]), it is unreasonable to assume that operators would expend resources keeping a fiber lit at orders of magnitude more capacity than is historically present on said link. In light

of this, ONSET assumes that every physical link in the network has one or more vacant channels (*i.e.*, frequency space that is not allocated to general traffic).

The second assumption is that all fiber spans carry signals only between nodes adjacent to the span. This assumption comes from applying ideas about IP-layer connectivity to optical networks. However, ROADMs and Optical Cross Connects (OXC) break this assumption, and allow for much more complex and diverse connectivity. Network designers often must make tough choices about where optical paths should be established based on trade-offs between IP demand, signal quality, and available channel space. Furthermore, these optical paths are prone to physical impairments (*e.g.*, BER/OSNR/Q Factor) and changes in signal quality can prompt re-routing of optical channels through the network only if space is available for the channel on an alternate path. In short, the optical layer is a heterogeneous collection of circuits between diverse endpoints; to simply saturate every link at maximum capacity oversimplifies this state of affairs and prohibits the cross-layer coordination of resources that we envision to be a powerful new tool for DDoS defense. We posit that bringing *optical awareness* (*i.e.*, connecting the optical layer to DDoS defenses in higher layers) leads to two key opportunities in combating present and future DDoS attacks.

2.1 Better performance for normal traffic in large-scale attacks

The first opportunity of optical-aware defense is to protect the performance of benign traffic by isolating it from suspicious traffic. Consider the scenario shown in Figure 1. Suppose that the two transit links, (RO1, RO2) and (RO2, RO3), each has 1 Tbps capacity and 10 ms latency, given by ten 100 Gbps optical channels, and that the average utilization of each link by benign users is 1/3. Then, 2/3 or 1.3 Tbps of capacity is left available for DDoS defence, flash crowds, link failures, *etc.* Also, since there is only one path from R1 to R3, the latency is ~20 ms. Now suppose that a DDoS attack occurs, and the scrubbing data center at R2 is equipped to filter all of the traffic, but that legitimate traffic incurs an additional latency of 200 ms. Now the latency for legitimate senders during an attack is 220 ms. We model the network cost by proxy using the sum of capacity for links in the network, therefore the network cost is 2.




Figure 1: Optical-unaware defense.

Using ONSET, if we physically separate trusted and suspicious traffic, we can gain performance and cost benefits. Recall that 2/3 of the link capacities are unused in general. Traffic is distributed among the ten optical channels evenly. Given that we know the typical usage of the network by legitimate users, and that we can detect an attack, suppose we can identify half of the flows on arrival as trusted. The capacity of trusted traffic then, is 1/3 legitimate users \times 1/2 trusted users \times 1 Tbps, or 167 Gbps of legitimate and trusted users. Now, we can send these flows on two express optical

channel from RO1 to RO3 with a latency of 10 ms. The performance benefit for trusted flows is 22× reduction in latency, and for all legitimate flows, the benefit is the weighted average of the trusted and suspicious/benign traffic, 4.5× reduction in latency ($1/6 \times 220/10 + 5/6 \times 220/220$). In order for a static IP topology to guarantee such performance, the physical layer connection between R01 and R03 needs to be available in perpetuity, adding an additional network cost of 1 Tbps. Therefore, ONSET achieves performance benefits while reducing network costs by 33%.




Figure 2: ONSET DDoS defense.

2.2 New capabilities for advanced/future infrastructure attacks

Beyond traffic isolation, the second opportunity of optical-aware defense is to dynamically reallocate capacities for under-attacked links. This could bring dual benefits, reducing the application latency and increasing the attacker’s cost.

Consider the setup shown in Figure 3. Suppose that all links have 1 Tbps of capacity, given by 10 optical channels and that link utilization by legitimate users is 1/3, or ~333 Gbps. In this case, 1.3 Tbps is reserved in anticipation of potentially performance-degrading events (e.g., DDoS, flash crowd, link failure, etc.). Again, we model the network cost via the sum of capacities for all links. Therefore, the network cost is 2. In this case, a link flooding attack needs 1M bots each sending 1 Mbps of traffic to render the link from RO1 to RO2 useless for legitimate users. (The attacker does not know about the legitimate user demand for the link, thus saturates the entire capacity.)




Figure 3: Optical-unaware defense for adaptive reconnaissance + flooding.

Using ONSET (Figure 4), we can both *scale the attacker’s cost* and *increase performance* using three distinct strategies:

(1) **Dynamic capacity expansion.** Recall, we know that link utilization by legitimate senders is 1/3 of the link’s capacity, which means there are six out of ten optical channels that can be re-provisioned away from link (RO1, RO3). If we allocate them onto (RO1, RO2), then the attack cost is scaled to 1.6M bots, a 60% increase. This new capacity can be used by Spiffy [31] to perform temporary bandwidth expansion (TBE) and identify malicious traffic more quickly. The cost for a static topology to provide the same capability

is 3.2, as it must drive the attack cost to 1.6M bots on both links simultaneously.

(2) **Dynamic capacity while providing an express route for trusted traffic.** Another defense strategy is to allocate a fraction (e.g., half) of the 6 channels between links (RO1, RO2), and (RO2, RO3). The new link from RO2 to RO3 provides an express-route for trusted traffic from R3 to R2, and the capacity given to (RO1, RO2) scales the attack cost while again providing more physical capacity for identifying malicious senders. In this case, the attack cost is increased to 1.3M bots, and performance is improved for a fraction of legitimate senders coming from R2. For a static topology to provide this benefit, it would need 1.3 Tbps provisioned between all RO nodes; this gives a total cost of 3.9.




Figure 4: ONSET defense for adaptive reconnaissance + link flooding.

(3) **Disrupting adaptive network reconnaissance.** To coordinate a link flooding attack, such as Crossfire [32], the attacker must infer the logical topology of the network. The attacker then places bots in the network and orchestrates communication between them such that they all use the targeted link. Now, if we allocate capacity onto a new logical link (e.g., RO2, RO3), this will change the behavior of the bot’s traffic. The attacker will have to restart their network reconnaissance and restart the attack. After they reconfigure the botnet, we can again reconfigure the logical topology, and force the attacker to start all over again. A static topology cannot mimic this effect.

3 ONSET VISION AND CHALLENGES

In this section, we describe the high-level vision of ONSET (§ 3.1), followed by the challenges (§ 3.2). Then, we model the potential benefits for performance during terabit attacks in § 3.3 with(out) a system like ONSET. Finally, we outline the design of our ongoing work in building ONSET (§ 3.4).

3.1 Vision

We envision ONSET: **O**ptics-enabled **I**n-**N**etwork **d**efen**S**e for **E**xtr**E**m**E** Terabit DDoS attacks. ONSET is a system to defend new-age and future extreme terabit DDoS attacks by leveraging optical layer programmability. Figure 5 depicts ONSET in contrast to traditional DDoS defenses. At the core of ONSET are three key insights: (a) isolate and steer suspicious traffic via opportunistic reconfiguration of wavelengths, ensuring the performance of benign traffic in the face of ongoing attacks; (b) provide coordination across relevant optical components, switch data planes, and network controllers via vendor-agnostic APIs; and (c) enhance the capacity of the network paths to alleviate attack-induced congestion in a performance-aware manner.

3.2 Challenges and Requirements

Realizing the above vision raises fundamental challenges—spanning modeling, systems, and networking aspects—which we explain below.

1. Gains of ONSET. Given the volume and variety of new-age and future DDoS attacks (distinguishable vs. indistinguishable, direct vs. indirect, and variable vs. fixed rate), what are the performance gains of ONSET in the face of larger scale legacy attacks? Answering this question is challenging due to paucity of analytical models for optical-enabled defenses for diverse larger-scale and advanced attack scenarios.

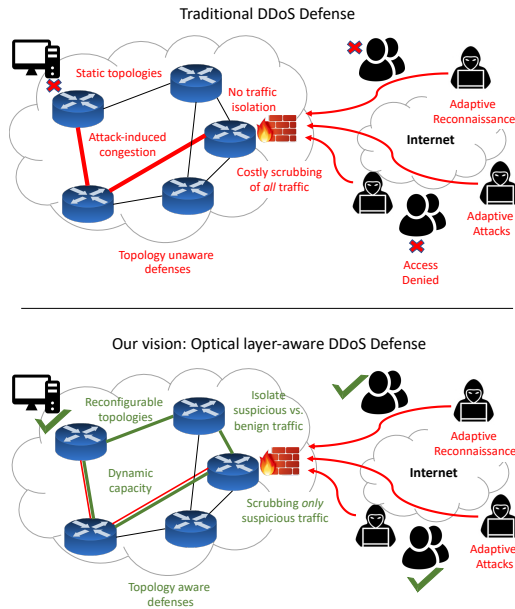


Figure 5: ONSET vision and opportunities.

2. Detection and reconfiguration in ONSET. How do we identify malicious vs. benign traffic at line rate and rapidly reconfigure the wavelengths and switches in ONSET? Addressing this is challenging due to a lack of lightweight detection capabilities at the data plane and unnecessary reconfiguration delays imposed by state-of-the-art optical gear.

3. Coordination and Management in ONSET. Further, how do we communicate the detected attacks and trigger new wavelengths on demand in a coordinated, performance-aware manner? While industry efforts (e.g., OpenConfig [41], Open Disaggregated Transport Network, or ODTN [2]) are working on packet-optical network coordination via open APIs, the scale and dynamism of new-age and future attacks mandate rethinking of these open APIs, and building resource management and orchestration modules to make ONSET practical.

3.3 Modeling Benefits of ONSET

WAN backbones are multi-million dollar assets for which the development and testing of new models and frameworks impose prohibitive start-up costs. While simulation frameworks exist for evaluating large packet-switched networks (e.g., Mininet [36]), these

do not allow for optical reconfigurability that a system like ONSET provides. The lack of robust and widely tested cross-layer (optical/packet) simulators is a new avenue of research for modeling and simulation of security applications. We take steps towards modeling our proposed defense and the performance benefit granted thereby in this section.

To quantify the benefits of ONSET-like systems, we consider the enterprise WAN modeled as a multi-graph $G = (V, E)$. V is a set of routers and switches and E is a multi-set of ordered pairs, i.e., $E = \{(e, c)\}$. Let e be an un-ordered pair of switches, $e = \{x, y\} : x, y \in V$, which represents a link (wavelength or lambda) between switches, and $c \in C$ be the capacity of the wavelength. These lambdas are *reconfigurable*, in the sense that they may be assigned on an edge, or any edge adjacent to a node for which the lambda is incident. In our model, flows (F) are denoted by a source and a destination address (src, dst). We assume that the flows originate and terminate outside of G , and categorize them into four subsets. Flows can be Attack (A), Suspicious (S), Not Suspicious (also known as Trusted) ($\neg S$), and Benign (B). Note, however, B is not necessarily trusted traffic ($\neg S$). The four classes are related as follows. All flows are either Suspicious or Trusted; $S \cup \neg S = F$ & $S \cap \neg S = \emptyset$. Attack traffic is a subset of Suspicious traffic; $A \subseteq S$. Trusted traffic is a subset of Benign traffic; $\neg S \subseteq B$. Benign traffic is $(S - A) \cup \neg S$.

We leverage this model to construct a system that increases performance for trusted traffic. In a traditional scrubbing-based solution, anomalies in traffic patterns trigger an alarm when voluminous traffic that is bound for a targeted client enters the network. After the attack is detected, the network reroutes all traffic bound for the target through a fixed set of hardware scrubbing appliances. With ONSET, however, we can achieve wavelength isolation of traffic, diverting only suspicious traffic S through the scrubber and forwarding $\neg S$ directly to the client. For instance, we can achieve this by switching $\neg S$ traffic to an alternate lambda *before* it can enter the data center. Then, a ROADM can be triggered to route the $\neg S$ wavelength directly to the destination so that the scrubber only processes suspicious traffic.

Preliminary Results. Let T Gbps of traffic enter the system. Let \hat{T} be the volume of traffic we forward to the client. If we scrub 100% of the ingress traffic, then the rate of traffic we forward to the client is limited by the bandwidth of the scrubbing appliances at the data center, T^D . However, if we can leverage reconfigurable lambdas to send trusted traffic, T^{-S} around the scrubbers, then the total bandwidth to the client is augmented with the trusted traffic (i.e., $\hat{T} = T^D + T^{-S}$).

To see this concretely, suppose a server is the target of a DDoS attack. As the volume of attack traffic increases, the volume of traffic that the server can respond to decreases as shown in Figure 6 (A). Now, suppose the victim has subscribed to a third party scrubbing service, and can detect the attack and re-route all traffic to the scrubber. This scrubbing service has a fixed capacity of 4 Tbps and is 100% effective in removing malicious traffic. However, after the attack volume exceeds the scrubber’s bandwidth, the total throughput for the victim decreases as seen in Figure 6 (B). If the attacker’s goal is to reduce the rate for legitimate senders to 300 Gbps, then they need to scale the attack from 2 Tbps to 8 Tbps—a 4× increase in cost. We assume that a fixed proportion of legitimate traffic can




Figure 6: (A) Baseline throughput for legitimate senders and malicious attackers when no defense is deployed. (B) Throughput for legitimate senders when using a 4 Tbps scrubber, with a baseline for reference of improvement. (C) Throughput for legitimate senders when 40% of good traffic is trusted. The ONSET line is the sum of the trusted line and scrubber line below. (D) ONSET’s throughput when 80% of traffic is trusted.

be trusted, regardless of the attack strength. For example, if the historical utilization of the service under attack is 700 Gbps, then a fraction of the senders who make up 700 Gbps of demand are labeled as trusted. With trusted traffic prioritized, we can forward it to the victim without involving a scrubber. Simultaneously, we deliver all other suspicious traffic to the scrubber.

Figures 6 (C) and (D) illustrate the throughput for different classes of traffic (trusted, suspicious, malicious, and total) when 40% or 80% of benign traffic is from trusted sources. We can see from the figure that goodput of the network (*i.e.*, the sum of throughput by trusted traffic and the scrubbers) asymptotically approaches the volume of trusted traffic as the strength of the attack grows. We also notice that the throughput for data leaving the scrubber approaches zero as the strength of the attack increases. We argue that physically separating traffic on distinct wavelengths, and only sending suspicious traffic to the scrubber, increases the quality of service for the victim’s network during a DDoS attack.

Furthermore, as the scale of the attack increases, so does the benefit in providing an optical layer-aware defense. When the attack strength reaches 10 Tbps in the baseline scenario, the throughput for legitimate senders falls to 65 Gbps, ~9% of its original strength. The scrubber alone, helped to keep throughput up to 260 Gbps, or ~37% of the full strength. If we forward 40% of the legitimate traffic as trusted, the aggregate throughput increases from 260 Gbps to 440 Gbps, or 62% of 700 Gbps—an improvement of 25% percent. Finally, if the network can identify 80% of traffic as trusted, then throughput for legitimate senders is ~620 Gbps, or 88%—an improvement of 51%. These early results show that ONSET can help increase throughput from 25% to 51% over the scrubber in this scenario.

3.4 Design Challenges

Motivated by the results above, in this section, we present the design of ONSET leaving its evaluation to future work.

ONSET uses optical-layer programmability to provide new defenses for DDoS attacks as illustrated in Figure 7. What makes ONSET unique is its ability to change the data plane of the network at the physical layer in response to an attack. ONSET defense applications are written using cross-layer APIs and configure the data plane’s optical and packet-switched network. Open protocols are used to facilitate these updates and are processed through a central controller with interfaces for managing the packet-switched

and optical network substrates. In this section, we describe the architecture in greater detail from the bottom up.

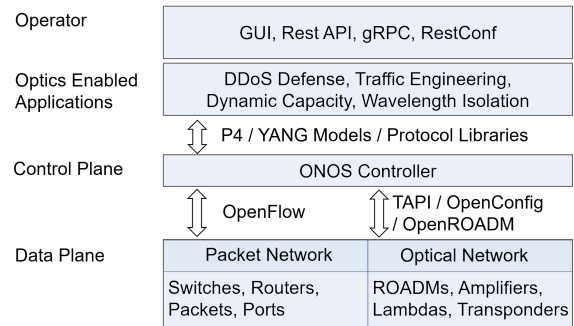


Figure 7: ONSET Architecture.

Data Plane: The data plane is built on the physical hardware that makes up the network, including all packet and optical hardware. Research on leveraging SDN to manipulate the packet-switched network for security purposes is vast, however optical-layer defense strategies are non-existent today; the reason being that the optical layer has primarily been inaccessible to higher-layer applications. Recent efforts by the Open Networking Foundation and other consortia of operators are changing this through the development of white-box data models for optical-layer hardware and open-source implementations.

Control Plane: Open Network Operating System (ONOS), is a controller that was built to bring greater flexibility to data plane management. It supports OpenFlow for packet-switch reconfiguration as well as several different interfaces for optical layer reconfiguration including Transport API (TAPI), OpenROADM, and OpenConfig. These optical layer interfaces have their own strengths and weaknesses. For example, TAPI is designed for network-wide configuration and thus keeps an up-to-date representation of the topology in the controller’s domain, however individual devices of the topology are out of TAPIs scope. OpenConfig and OpenROADM have greater support for device-level monitoring and configuration. For example, OpenConfig allows manual amplifier-gain adjustment and even wavelength-power adjustments on line-terminating equipment. OpenROADM’s support is somewhere between TAPI and OpenConfig; it is concerned with curating a holistic view of the

network while enabling operators to configure lower-level details of transport intents (e.g., routing-wavelength assignments).

Optics Enabled Applications: ONOS is extensible and therefore allows integration with user-written applications. Applications can be written in P4 (e.g., to relate to a switch how it should forward traffic given different instances of an underlying optical topology), as YANG data models, and in terms of protocol libraries. The vision of ONSET is to leverage this extensibility to create a new suite of defense applications that realize the potential for reconfigurable optics. Some example applications include—but are not limited to—DDoS defense, dynamic capacity, wavelength isolation of traffic, and traffic engineering.

Operator Interface: ONSET supports all of the ONOS operator interfaces, including GUI, Rest API, gRPC, RestConf, and others. Operators can push new applications to ONSET via these interfaces, and ensure consistent and ACID-compliant changes to the control infrastructure and data plane.

Hitless Transitions: ONSET should ideally guarantee that traffic traveling across a path that is being reconfigured arrives at its ultimate destination. The time taken for a reconfiguration event must be as short as possible for hitless transitions to be feasible. Furthermore, the network can rely on alternative paths to forward traffic while migrating lambdas from one set of optical links to another. In fact, such systems for delivering traffic in dynamic optical network settings have been demonstrated in data center networks [22, 40]. As the time for reconfiguration is brought down in WAN networks, these solutions could be adapted for ONSET.

4 RELATED WORK

SDN + Optics. The use of reconfigurable wavelengths in ONSET is similar to prior work on software-defined elastic optical networks (SD-EONs); e.g., see [11, 12, 15, 38, 49]. The main goal of these efforts is to realize cross-domain lightpath provisioning with cooperative games [15, 49] and not flexible DDoS defenses. On the contrary, ONSET's goal is to allocate wavelengths to isolate and mitigate DDoS attacks. Moreover, the tasks described in this proposal focus on bringing optical awareness to today's DDoS defenses. SDN-based infrastructure provisioning, in the context of datacenter and WAN settings, include B4 [28], SWAN [24], Owan [30], and others [14, 23, 33, 45, 51], each of which aims at improving the utilization of inter-data center and WANs. A survey of related efforts is available here [6, 50]. None of these efforts focus on DDoS.

SDN + Security. Prior work has shown the benefits of SDN and NFV to make network security more flexible. Ethane [10] uses a centralized controller controls switches at critical points to authorized traffic. FlowGuard [26] resolves interfering ACL policies from firewalls in the SDN network. Other efforts such as FRESKO [44] implement detection and mitigation modules in the SDN controller. However, the controller becomes a critical bottleneck for scalability and requires reimplementing functionality that is commonly available in security middleboxes. OFX [48] and Kinetic [34] focus on security policies for networks composed of switches. PBS [25] addresses security challenges induced by BYOD using SDN. Other efforts have also considered domains such as IoT security. These efforts do not focus on DDoS attacks which is the focus of this proposal. More closely related to our work are efforts on looking at

SDN-NFV solutions for DDoS mitigation (e.g., [19, 53]). At a high level, these efforts are complementary techniques that orchestrate well-known defenses on VMs and switches and do not consider the novel dimensions of optical/topology reconfiguration which is the focus of this work. While other work [31, 47] use route reconfiguration for DDoS defense, they do so at the IP layer and do not consider the full spectrum of recent optical programming capabilities.

5 FUTURE OUTLOOK

An approach like ONSET opens up a number of interesting problems at the intersection of optical, security, and networking communities, which we outline below.

Feasibility of ONSET for Diverse DDoS Attacks. Apart from the direct DDoS presented, grand challenges lie ahead in modeling and evaluating the gains of ONSET for combating other types of DDoS attacks. In particular, the feasibility of ONSET in defending (in)distinguishable, fixed vs. variable rate, volume-based, and protocol-conforming attacks calls for further research involving optical and security communities. To this end, we are working to build an accurate modeling and simulation platform for ONSET. The goal of this platform is to extend existing simulators (e.g., ONOS Packet/Optical [3], ODTN-Emulator [2]) and enable experiments to support ONSET security applications.

Prototyping ONSET. The success of ONSET depends on building a prototype based on the design outlined above and its rigorous evaluation in the face of diverse attack scenarios. This, in turn, calls for collaborations among optical, networking, and security researchers. Furthermore, the heterogeneity, scale, and dynamism of modern DDoS attacks require new testing frameworks and capabilities for ONSET to operate effectively against the growing DDoS landscape.

ONSET for Advanced Cyber Attacks. While the goal of this paper is to make a case for optical layer-aware DDoS defense, we believe that the notion of optical layer awareness is beneficial for a broader class of cyber attacks. First, reconnaissance is an on-going problem since the topology can be mapped as shown by Achleitner et al. [4]. By keeping the performance and network objectives in mind, we believe ONSET can arbitrarily change wavelengths to effectively combat reconnaissance by providing cyber deception. In addition, complementary to NetHide [39], we believe that ONSET can be used to combat targeted attacks by dynamically altering the underlying wavelengths and, hence, topologies.

Acknowledgments

We thank the anonymous reviewers for their feedback. This work was supported by UO faculty research award, Ripple fellowship, and in part by NSF awards CNS-1513764, CNS-1700521, and CNS-1850297. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NSF, UO, or Ripple.

REFERENCES

- [1] DDoS Attacks Up By 84% in Q1. <https://www.cybersecurityintelligence.com/blog/ddos-attacks-up-by-84-in-q1-4346.html>.
- [2] Open disaggregated transport network. <https://www.opennetworking.org/odtn/>.
- [3] Open network operating system (onos). <https://www.opennetworking.org/onos/>.

- [4] S. Achleitner, T. La Porta, P. McDaniel, S. Sugrim, S. V. Krishnamurthy, and R. Chadha. Cyber deception: Virtual networks to defend insider reconnaissance. In *Proceedings of the 8th ACM CCS international workshop on managing insider security threats*, pages 57–68. ACM, 2016.
- [5] Akamai. Akamai security solutions. <https://www.akamai.com/us/en/products/security/>, 2019.
- [6] R. Alvizu, G. Maier, N. Kukreja, A. Pattavina, R. Morro, A. Capello, and C. Cavazoni. Comprehensive survey on t-sdn: Software-defined networking for transport networks. *IEEE Communications Surveys & Tutorials*, 2017.
- [7] AWS. Aws shield: Managed ddos protection. <https://aws.amazon.com/shield/>, 2019.
- [8] F. Baker and P. Savola. Ingress filtering for multihomed networks. Technical report, BCP 84, RFC 3704, March, 2004.
- [9] E. Bursztein. Inside mirai the infamous iot botnet: A retrospective analysis. <https://elie.net/blog/security/inside-mirai-the-infamous-iot-botnet-a-retrospective-analysis/>, dec 2018.
- [10] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. McKeown, and S. Shenker. Ethane: Taking control of the enterprise. In *ACM SIGCOMM Computer Communication Review*, volume 37, pages 1–12. ACM, 2007.
- [11] A. Castro, L. Gifre, C. Chen, J. Yin, Z. Zhu, L. Velasco, , and S.-J. B. Yoo. Experimental demonstration of brokered orchestration for end-to-end service provisioning and interoperability across heterogeneous multi-operator (multi-as) optical networks. *IEEE European Conference on Optical Communication*, 2015.
- [12] A. Castro, L. Velasco, L. Gifre, C. Chen, J. Yin, Z. Zhu, R. Proietti, and S.-J. B. Yoo. Brokered orchestration for end-to-end service provisioning across heterogeneous multi-operator (multi-as) optical networks. *IEEE Journal of Lightwave Technology*, 2016.
- [13] CenturyLink. Centurylink ddos mitigation. <http://www.centurylink.com/asset/business/enterprise/brochure/ddos-mitigation.pdf>, 2019.
- [14] M. Channegowda, R. Nejabati, and D. Simeonidou. Software-defined optical networks technology and infrastructure: Enabling software-defined optical network operations. *Journal of Optical Communications and Networking*, 2013.
- [15] X. Chen, J. Yin, C. Chen, Z. Zhu, A. Casales, and S. Yoo. Multi-broker based market-driven service provisioning in multi-domain sd-eons in noncooperative game scenarios. 2015.
- [16] Cloudflare. Advanced ddos attack protection. <https://www.cloudflare.com/ddos/>, 2019.
- [17] R. Durairajan, P. Barford, J. Sommers, and W. Willinger. Greyfiber: A system for providing flexible access to wide-area connectivity. *arXiv preprint arXiv:1807.05242*, 2018.
- [18] A. D. Ellis, J. Zhao, and D. Cotter. Approaching the non-linear shannon limit. *Journal of Lightwave Technology*, 28(4):423–433, 2009.
- [19] S. K. Fayaz, Y. Tobioka, V. Sekar, and M. Bailey. Flexible and Elastic DDoS Defense Using Bohatei . In *Proc. USENIX Security*, 2015.
- [20] P. Ferguson and D. Senie. Network ingress filtering: Defeating denial of service attacks which employ ip source address spoofing. RFC 2827, RFC Editor, May 2000.
- [21] M. Gaiser. How much monetary damage was done during the oct 21, 2016 ddos of dyndns? <https://www.quora.com/How-much-monetary-damage-was-done-during-the-Oct-21-2016-DDOS-of-DynDNS>, Oct 2016.
- [22] M. Ghobadi, R. Mahajan, A. Phanishayee, N. Devanur, J. Kulkarni, G. Ranade, P.-A. Blanche, H. Rastegarfar, M. Glick, and D. Kilper. Projector: Agile reconfigurable data center interconnect. In *Proceedings of the 2016 ACM SIGCOMM Conference*, pages 216–229, 2016.
- [23] A. Giorgetti, F. Paolucci, F. Cugini, and P. Castoldi. Dynamic restoration with gmpls and sdn control plane in elastic optical networks. *Journal of Optical Communications and Networking*, 2015.
- [24] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer. Achieving High Utilization with Software-driven WAN. In *ACM SIGCOMM*, 2013.
- [25] S. Hong, R. Baykov, L. Xu, S. Nadimpalli, and G. Gu. Towards sdn-defined programmable byod (bring your own device) security. In *NDSS*, 2016.
- [26] H. Hu, W. Han, G.-J. Ahn, and Z. Zhao. Flowguard: building robust firewalls for software-defined networks. In *Proceedings of the third workshop on Hot topics in software defined networking*, pages 97–102. ACM, 2014.
- [27] R. Ikhsan, R. F. Syahputra, et al. Performance control of semiconductor optical amplifier and fiber raman amplifier in communication system. In *2018 19th IEEE/ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD)*, pages 32–36. IEEE, 2018.
- [28] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, et al. B4: Experience with a Globally-deployed Software Defined WAN. *ACM SIGCOMM*, 2013.
- [29] C. Jin, H. Wang, and K. G. Shin. Hop-count filtering: an effective defense against spoofed ddos traffic. In *Proceedings of the 10th ACM conference on Computer and communications security*, pages 30–41. ACM, 2003.
- [30] X. Jin, Y. Li, D. Wei, S. Li, J. Gao, L. Xu, G. Li, W. Xu, and J. Rexford. Optimizing Bulk Transfers with Software-Defined Optical WAN. *ACM SIGCOMM*, 2016.
- [31] M. S. Kang, V. D. Gligor, and V. Sekar. Spiffy: Inducing cost-detectability tradeoffs for persistent link-flooding attacks. In *NDSS*, 2016.
- [32] M. S. Kang, S. B. Lee, and V. D. Gligor. The crossfire attack. In *2013 IEEE Symposium on Security and Privacy*, pages 127–141. IEEE, 2013.
- [33] D. C. Kilper and Y. Li. Optical physical layer sdn: Enabling physical layer programmability through open control systems. In *Optical Fiber Communications Conference and Exhibition (OFC)*, 2017, 2017.
- [34] H. Kim, J. Reich, A. Gupta, M. Shahbaz, N. Feamster, and R. Clark. Kinetic: Verifiable dynamic network control. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, pages 59–72, Oakland, CA, May 2015. USENIX Association.
- [35] M. N. Kumar, P. Sujatha, V. Kalva, R. Nagori, A. K. Katukojwala, and M. Kumar. Mitigating economic denial of sustainability (edos) in cloud computing using in-cloud scrubber service. In *2012 Fourth International Conference on Computational Intelligence and Communication Networks*, pages 535–539. IEEE, 2012.
- [36] B. Lantz, B. Heller, and N. McKeown. A Network in a Laptop: Rapid Prototyping for Software-defined Networks. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, 2010.
- [37] M. Li, D. Zaccarin, and C. Barnard. Reconfigurable optical add-drop multiplexer, Feb. 27 2007. US Patent 7,184,666.
- [38] D. Marconett and S. Yoo. Flowbroker: Market-driven multi-domain sdn with heterogeneous brokers. In *Optical Fiber Communication Conference*, pages Th2A–36. Optical Society of America, 2015.
- [39] R. Meier, P. Tsankov, V. Lenders, L. Vanbever, and M. Vechev. Nethide: secure and practical network topology obfuscation. In *27th USENIX Security Symposium*, pages 693–709, 2018.
- [40] W. M. Mellette, R. Das, Y. Guo, R. McGuinness, A. C. Snoeren, and G. Porter. Expanding across time to deliver bandwidth efficiency and low latency. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI '20)*, pages 1–18, 2020.
- [41] OpenConfig. Vendor-neutral, model-driven network management designed by users. <http://www.openconfig.net/>, 2016.
- [42] G. Porter, R. Strong, N. Farrington, A. Forencich, P. Chen-Sun, T. Rosing, Y. Fainman, G. Papen, and A. Vahdat. *Integrating microsecond circuit switching into the data center*, volume 43. ACM, 2013.
- [43] J. Russell. The world's largest ddos attack took github offline for fewer than 10 minutes. <https://techcrunch.com/2018/03/02/the-worlds-largest-ddos-attack-took-github-offline-for-less-than-tens-minutes/>, 2018.
- [44] S. Shin, P. A. Porras, V. Yegneswaran, M. W. Fong, G. Gu, and M. Tyson. Fresco: Modular composable security services for software-defined networks. In *NDSS*, 2013.
- [45] R. Singh, M. Ghobadi, K.-T. Förster, M. Filer, and P. Gill. Run, walk, crawl: Towards dynamic link capacities. In *Proceedings of the 16th ACM HotNets 2017*, 2017.
- [46] A. Singla, A. Singh, K. Ramachandran, L. Xu, and Y. Zhang. Proteus: a topology malleable data center network. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, page 8. ACM, 2010.
- [47] J. M. Smith and M. Schuchard. Routing around congestion: Defeating ddos attacks and adverse network conditions via reactive bgp routing. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 599–617. IEEE, 2018.
- [48] J. Sonchack, J. M. Smith, A. J. Aviv, and E. Keller. Enabling practical software-defined networking security applications with ofx. In *NDSS*, 2016.
- [49] L. Sun, X. Chen, and Z. Zhu. Multibroker-based service provisioning in multidomain sd-eons: Why and how should the brokers cooperate with each other? *IEEE Journal of Lightwave Technology*, 2017.
- [50] A. S. Thyagaturu, A. Mercian, M. P. McGarry, M. Reisslein, and W. Kellerer. Software defined optical networks (sdons): A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 2016.
- [51] Y. Xiong, Y. Li, B. Zhou, R. Wang, and G. N. Rouskas. Sdn enabled restoration with triggered precomputation in elastic optical inter-datacenter networks. *Journal of Optical Communications and Networking*, 2018.
- [52] A. Yaar, A. Perrig, and D. Song. Stackpi: New packet marking and filtering mechanisms for ddos and ip spoofing defense. *IEEE Journal on Selected Areas in Communications*, 24(10):1853–1863, 2006.
- [53] M. Zhang, G.-Y. Li, S. Wang, C. Liu, A. Chen, H. Hu, G. Gu, Q. Li, M. Xu, and J. Wu. Poseidon: Mitigating volumetric ddos attacks with programmable switches. 2020.