Aggregate-Based Congestion Control
for Pulse-Wave DDoS Defense

Albert Gran Alcoz
ETH Zürich

Martin Strohmeier
Armasuisse

Vincent Lenders
Armasuisse

Laurent Vanbever
ETH Zürich

ABSTRACT

Pulse-wave DDoS attacks are a new type of volumetric attack formed by short, high-rate traffic pulses. Such attacks target the Achilles' heel of state-of-the-art DDoS defenses: their reaction time. By continuously adapting their attack vectors, pulse-wave attacks manage to render existing defenses ineffective.

In this paper, we leverage programmable switches to build an in-network DDoS defense effective against pulse-wave attacks. To do so, we revisit Aggregate-based Congestion Control (ACC): a mechanism proposed two decades ago to manage congestion events caused by high-bandwidth traffic aggregates. While ACC proved efficient in inferring and controlling DDoS attacks, it cannot keep up with the speed requirements of pulse-wave attacks.

We propose *ACC-Turbo*, a renewed version of ACC that infers attack patterns by applying online-clustering techniques in the network and mitigates them by leveraging programmable packet scheduling. By doing so, *ACC-Turbo* identifies attacks at line rate and in real-time, and rate-limits attack traffic on a per-packet basis.

We fully implement *ACC-Turbo* in P4 and evaluate it on a wide range of attack scenarios. Our evaluation shows that *ACC-Turbo* autonomously identifies DDoS attack vectors in an unsupervised manner and rapidly mitigates pulse-wave DDoS attacks. We also show that *ACC-Turbo* runs on existing hardware (Intel Tofino).

CCS CONCEPTS

• **Networks → Network security**; **Denial-of-service attacks**;

KEYWORDS

Network Security, DDoS, Pulse-Wave DDoS, ACC, Aggregate-Based Congestion Control, Programmable Scheduling, Network Defenses

1 INTRODUCTION

Pulse-wave DDoS attacks have recently managed to take down critical network infrastructure while causing enormous financial and reputational damages [2, 4, 49, 57]. In contrast to conventional DDoS attacks, which grow steadily and persist longer in time, pulse-wave DDoS attacks consist of high-rate short-lived bursts. Each burst typically leverages a different attack vector (e.g., NTP, DNS, Memcached) and can reach hundreds of Gbps [38, 50, 52].

The threat in pulse-wave attacks resides in that they target the Achilles' heel of existing DDoS mitigation systems: their reaction time. Most in-network DDoS defenses today (both in research and production) rely on some sort of offline facility that can either directly scrub traffic [7, 19, 23, 34, 42, 58], orchestrate a routing-based defense [46, 51], or deploy an in-network pre-configured mitigation [35, 54]. The time required for an in-network defense to reach this external facility and deploy the corresponding defense can be in the order of seconds to minutes [52]. Pulse-wave attacks exploit this vulnerability by sending traffic pulses that force the DDoS defense to repeat this control loop over and over. By keeping the defenses in a constant transient state, pulse-wave attacks manage to make them ineffective. If the defenses rely on traffic redirection, pulse-wave attacks may even produce route flapping in the network [26].

Designing a pulse-wave defense is intricate. Like conventional-DDoS defenses, an ideal pulse-wave defense needs to be, first, *generic*, to identify a wide variety of attack vectors at different granularity [5, 17]. Generic techniques usually require unsupervision and incur the risk of misclassifying traffic. Thus, an ideal defense also needs to be *measured* in responding to attacks [47]. Most DDoS defenses fail at the first condition. For example, signature-based defenses [35, 54, 58] are not generic. They only cover a small set of attack vectors, and can not keep up with the constantly-growing list of new attacks [5, 17]. Similarly, most congestion-management tools (e.g., heavy-hitter detectors or active queue management) lack granularity: they only work at, e.g., the per-flow level or the whole-traffic level. Other DDoS defenses fail at satisfying the second condition. For example, drop- or routing-based defenses [20, 39] strongly degrade performance in case of misclassification.

Aggregate-Based Congestion Control (ACC) was proposed two decades ago, *satisfying the two* design conditions [36]. ACC is a canonical mechanism to reduce the impact of congestion caused by *generic* traffic aggregates, with a *measured* bandwidth control. At a high level, ACC is a feedback loop that iteratively: *(i)* infers the aggregates causing the congestion; before *(ii)* reducing their throughput to a reasonable level. To infer the aggregates, ACC clusters the headers of packets dropped by a Random Early Detection (RED) queuing discipline. ACC then rate-limits the inferred aggregates to keep the total traffic throughput below the link capacity.

While ACC is effective at inferring and controlling conventional DDoS attacks, it fails at mitigating pulse-wave attacks, as it cannot keep up with the required fast reaction times. The reason is two-fold: (i) ACC relies on *offline* inference and control mechanisms, which run on either a separate server or a control plane, and (ii) ACC leverages a threshold-based defense activation. By running *offline*, ACC suffers from slow reaction time, which opens the door to pulse-wave DDoS attacks. By relying on threshold-based activation, ACC either further slows down reaction time or increases the probability of false positives, which negatively impacts its performance. In our experiments, even with the best configuration, ACC drops ≈ 20% of benign traffic in the event of a pulse-wave attack (§2.2).

Our work. In this paper, we redesign ACC for pulse-wave DDoS defense. We propose *ACC-Turbo*: the *first* sub-second-reaction-time aggregate-based congestion control mechanism that mitigates pulse-wave DDoS attacks by running at line-rate on commodity hardware. *ACC-Turbo* strategically combines two key techniques: *online clustering* directly in the data plane to infer attacks, and *programmable scheduling* to mitigate them (cf. Table 1).

First, *ACC-Turbo* offloads the clustering process to the data plane and analyzes *all* the traffic at line rate (instead of just a sample). This allows *ACC-Turbo* to speed up reaction time substantially. Further, since data-plane processing does not impact traffic latency, *ACC-Turbo* runs the clustering algorithm *continuously*. This eliminates the need for a threshold-based activation, which can be vulnerable to pulse-wave attacks and opens the door to potential false positives. The *always-on* design enables *ACC-Turbo* to anticipate the congestion events, achieving yet-faster reaction time.

Second, *ACC-Turbo* uses programmable scheduling to deprioritize malicious traffic instead of dropping or rate-limiting it. Doing so has three advantages: (i) it accommodates fine-grained assessments which can adjust to the requirements of individual aggregates, increasing fairness; (ii) it adapts at the per-packet level, being able to react to traffic variations rapidly, achieving faster and more-accurate bandwidth allocations; and (iii) it only leads to hard drops under congestion, being transparent (and therefore safe) otherwise.

Evaluation. We implement *ACC-Turbo* in P4 [11] and run it on programmable switches (Intel Tofino). We show that *ACC-Turbo* effectively mitigates pulse-wave DDoS attacks in real-time (i.e., less than 1*s* reaction time), rapidly adapting to attack variations. We also compare *ACC-Turbo* with Jaqen [35], a state-of-the-art DDoS defense. We show that *ACC-Turbo* is *at least* 10× faster than Jaqen in mitigating a broader range of attacks while also being safer.

Contributions. We make the following contributions:

- An online-clustering approach to infer pulse-wave DDoS attacks at scale directly from the network (§4).
- A scheduling algorithm to mitigate pulse-wave attacks and minimize their impact on background traffic (§5).
- An implementation[1] of *ACC-Turbo* in Python/P4 (§6).
- A comprehensive evaluation showing *ACC-Turbo*'s practicality and its ability to run on hardware (§7, §8).
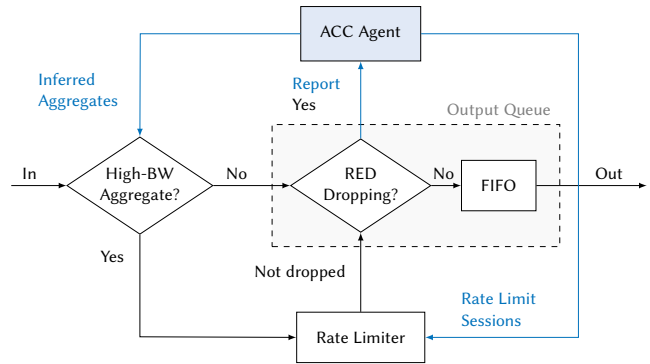
---

[1] Available at https://github.com/nsg-ethz/ACC-Turbo

Figure 1: ACC architecture [36].

2 BACKGROUND

In this section, we review the design of ACC (§2.1), as well as its limitations in mitigating pulse-wave DDoS attacks (§2.2).

2.1 Aggregate-based Congestion Control

Aggregate-based Congestion Control is a mechanism for inferring and controlling high-bandwidth aggregates that persistently overload a link in the network. In this paper, we focus on its local version, which runs on the switch that gives access to the congested link.[2]

Fig. 1 shows the architecture of an ACC-enabled switch. The core of ACC is composed of a RED module, implemented on top of a First-In First-Out (FIFO) queue. This module monitors the average queue size of the FIFO queue and drops packets probabilistically depending on its size. If the FIFO queue is almost empty, all incoming packets are accepted. As the queue grows, the probability of dropping an incoming packet increases. When the queue is full, the probability is at its maximum, and all incoming packets are dropped.

Whenever the RED module decides to drop a packet, it reports the header of the dropped packet to an ACC Agent. The ACC Agent periodically analyzes the packet headers of all dropped packets and tries to infer the traffic aggregates responsible for the congestion. When the ACC Agent identifies an aggregate, it creates a rate-limiting session to police the throughput of the aggregate. From that point onwards, all upcoming packets belonging to the aggregate will be rate-limited before being processed by the RED module.

Identifying congestion. The ACC Agent is activated when the output queue experiences sustained high congestion. This occurs when the drop rate in the output queue exceeds a pre-defined value, p_{high}, during a certain time period, K, which is also pre-defined.

Inferring aggregates. The ACC Agent infers traffic aggregates based *solely* on IP prefixes. At the high level, it extracts a list of either source or destination IP addresses that account for more than twice the mean number of packet drops and clusters them into 24-bit prefixes. To minimize collateral damage, ACC then walks down the prefix subtree for each of the aggregates, trying to obtain longer prefixes that still contain most of the packet drops.

---

[2] The original work also includes a "pushback" mechanism to extend the rate-limiting policies to upstream switches. This part is out of our scope.

	Activation	Inference	Control
ACC [36]	Threshold-based	Offline clustering on RED drops	Uniform rate limiting using estimated rates
ACC-Turbo	Always-on	Online clustering on all traffic	Programmable scheduling using exact rates

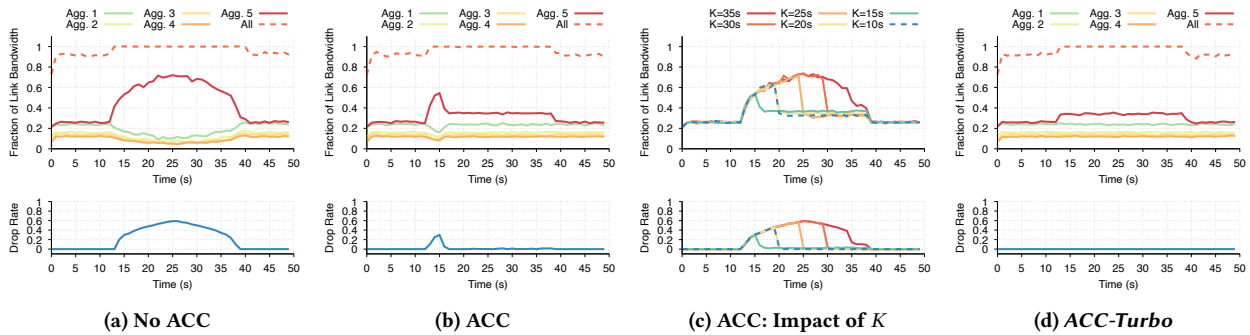Table 1: *ACC-Turbo* techniques vs. ACC.

Figure 2: Comparison between ACC and *ACC-Turbo* with the original papers' experiment [36].

Rate-limiting aggregates. For each aggregate inferred, the ACC Agent then computes the bandwidth to which it should be rate-limited. This limit is computed such that the drop rate at the output queue, gets below a pre-defined value, p_{target}. To that end, the ACC Agent, first, sorts the list of inferred aggregates by their number of drops (highest, first). Then, it computes the excess arrival rate at the output queue, R_{excess}, defined as the amount of traffic that should be dropped in order for the drop rate to get below p_{target}. Finally, the ACC Agent determines the minimum number of aggregates that should be rate-limited, $|\mathcal{A}|$, and the rate to which they should be limited, L, such that the total rate is reduced by R_{excess}:

$$\sum_{i=1}^{|\mathcal{A}|}(Aggregate[i].rate - L) = R_{excess}$$

With this design, ACC manages to be generic (inferring attacks agnostically to their characteristics) and measured (rate-limiting inferred attacks instead of just dropping them).

Experiment. We illustrate ACC's performance by using packet-level simulations (cf., §8 for more details about the setup). We reproduce ACC's original experiment [36]. It consists in scheduling five aggregates (1-5) over a bottleneck link using FIFO, and ACC, respectively. Aggregates 1-4 are constant-bit-rate flows. Aggregate 5 is a variable-rate flow which represents an attack, and starts increasing (resp. decreasing) its rate at $t = 13s$ (resp. $t = 25s$).

Fig. 2a and Fig. 2b show the bandwidth share across the five aggregates (top) and the drop rate at the output queue (bottom) when no protection (i.e., FIFO) and ACC are used, respectively. ACC is configured with $p_{high} = 0.1$, $K = 2s$, and $p_{target} = 0.05$. Appendix A details the rest of parameters. Without ACC, we see how the attack traffic captures most of the link bandwidth, degrading the performance of the other aggregates. With ACC, the attack is efficiently mitigated. Indeed, when the drop rate at the output queue exceeds p_{high}, the ACC Agent infers the attack and rate-limits it sufficiently to reduce its impact on background traffic.

Finally, we also see how the reaction time of ACC is $\approx 4s$. This is the time since the first attack packets arrive (at $t = 13s$), until the defense is deployed (at $t = 17s$). This time is mostly driven by the monitoring-window size, K. Indeed, for smaller K values, ACC checks more often whether p_{high} is exceeded, being able to identify faster whenever that occurs. A lower K, however, does not always imply a faster reaction time (cf. Fig. 2c). For example, $K = 10s$ achieves a slower reaction time than $K = 15s$. The reason is that when $K = 10s$, even though the threshold is checked first at $t = 10s$, it is not triggered until $t = 20s$, when p_{high} is reached.

2.2 Limitations of ACC

While ACC manages to defend against conventional DDoS attacks successfully, it is vulnerable to pulse-wave DDoS attacks. The reason is two-fold. First, ACC uses *offline* inference and control mechanisms, which run on either a separate server or a control plane [36]. Pulse-wave DDoS attacks are specially crafted to target the time required by such control planes to compute and deploy the right DDoS defense. They do so by sending short high-rate traffic pulses which morph over time. These pulses congest the link resources while the inference and control processes are running. By the time the inference and control processes manage to converge and ACC mitigates the attack, another pulse comes in, forcing the control loop to start the mitigation process all over again.

The second reason is that ACC relies on a threshold-based defense activation, which introduces a second vulnerability to pulse-wave attacks: If we configure a too small threshold (to speed up reaction time), the probability of false positives (i.e., benign traffic bursts identified as attacks) increases, as we prove in §8. In contrast, if we configure a too large threshold (to maximize accuracy), reaction time gets slower, opening the door to pulse-wave attacks. In ACC, false positives are especially concerning under attack, given that (i) ACC rate-limits all aggregates to the same amount, and (ii) its rate-limiting policies have long-lasting effects (cf. §2.1).
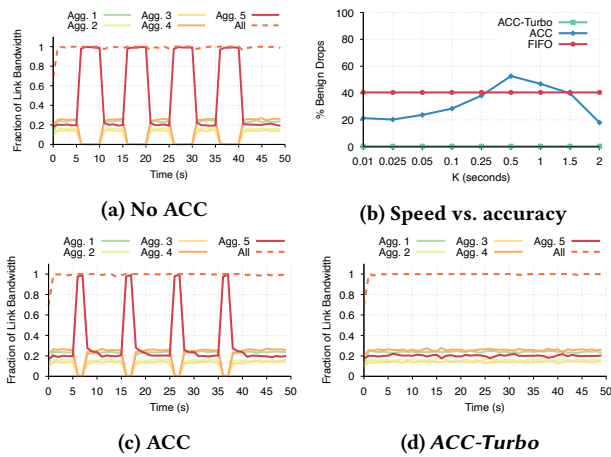
(a) No ACC

(b) Speed vs. accuracy

(c) ACC

(d) ACC-Turbo

Figure 3: Performance under morphing attack.

Example. In the following, we evaluate ACC's mitigation efficiency in the case of a pulse-wave attack composed of four vectors (starting at $5s$, $15s$, $25s$, and $35s$). For simplicity, we represent the four pulses as a single "attack" aggregate (i.e., aggregate 5). We leverage four constant-bit-rate flows as background traffic (i.e., aggregates 1-4), which transmit at $\approx$ the link capacity.

Fig. 3a and Fig. 3c illustrate the bandwidth share at the output link when FIFO and ACC (configured as in §2.1) are used. We see how ACC fails at mitigating the attack, only managing to defend the second half of attack pulses. By the time the monitoring window K is triggered, and p_{high} is reached, the pulse-wave attack has already managed to throttle benign traffic. Fig. 3b shows how reducing the size of the monitoring window, K, does not help. In fact, due to false positives, the performance also gets bounded for the smallest K values, with $\approx 20\%$ of benign traffic being dropped.

Given the incapacity of ACC to mitigate pulse-wave DDoS attacks, we propose *ACC-Turbo*. With an under-second reaction time, and disposing of the threshold-based activation, *ACC-Turbo* manages to successfully mitigate *both* conventional *and* pulse-wave DDoS attacks (cf., Fig. 2d and Fig. 3d, respectively).

3 OVERVIEW

In this section, we present the threat model (§3.1), and introduce an overview of *ACC-Turbo*'s design (§3.2).

3.1 Threat model

Attacker. The attacker's objective is to generate a pulse-wave DDoS attack (i.e., a series of short-duration high-rate traffic pulses) towards a critical link of the network in order to exhaust its capacity and prevent legitimate flows from using it. To that end, the attacker is free to generate *any* kind of attack traffic. For instance, the attacker can rely on: (i) a botnet of infected devices which directly floods traffic towards the link; (ii) reflection and amplification techniques, which send spoofed requests to open servers such that their responses cross the link; or even (iii) complex link-flooding attacks which exchange low-rate flows from numerous sources to numerous destinations such that they also cross the link [29].
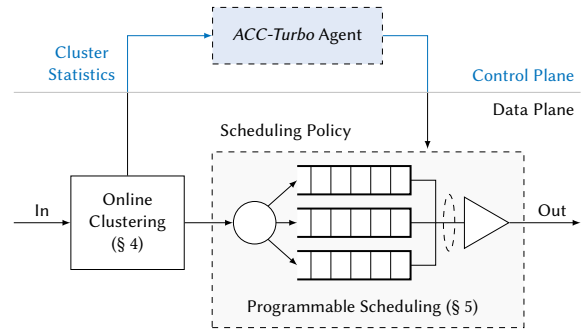

Figure 4: *ACC-Turbo* architecture.

System model. Same as ACC, *ACC-Turbo* runs on the switch that gives access to the critical link. This switch can be, e.g., an edge-router at an ISP or IXP, with a bigger input capacity than the output link's bandwidth. *ACC-Turbo* analyzes all traffic entering the switch (including attack traffic) and processes each packet individually. We consider that *ACC-Turbo* only looks at the packet headers, and it performs limited computations. (This is imposed by the fact that we want *ACC-Turbo* to run at line rate on programmable switches [11]).

3.2 *ACC-Turbo* design

ACC-Turbo is a switch-native aggregate-based congestion control mechanism that leverages programmable switches to mitigate pulse-wave DDoS attacks. It consists of an *online-clustering* module that runs entirely in the data plane and a *programmable-scheduling* module that runs in both the control plane and the data plane (see Fig. 4). With this hybrid design, *ACC-Turbo* balances the need for fast reaction (by running the inference process in the data plane) and accuracy (by dedicating most data-plane resources to inference while offloading the least time-sensitive parts to the control plane).

The online-clustering module extracts a set of features from the headers of arriving packets and uses them to cluster similar packets together. It runs in the data plane, processing *all* packets at line rate. It also runs *continuously*, regardless of whether there is congestion, eliminating the need for a threshold-based activation. Running continuously also allows *ACC-Turbo* to anticipate its inference decisions to the actual attack, speeding up reaction time.

The programmable-scheduling module involves both the control and the data plane. The control plane periodically polls information about the extracted clusters, including statistics about the clusters and their exact arrival rates. Then, it uses this information to assess the probability that each cluster contains attack traffic, and derives a scheduling policy for each cluster. The scheduling policies aim at deprioritizing clusters with a higher probability of being part of an attack. *ACC-Turbo* finally deploys these policies to the data plane and applies them to incoming packets mapped to each cluster.

By using programmable scheduling, *ACC-Turbo* adapts to traffic variations at a per-packet granularity and rapidly reacts to attack changes. Further, programmable scheduling enables *ACC-Turbo* to run continuously, even in the case of no attack. This is because programmable scheduling does not hurt traffic: it only drops packets in case of severe congestion (being transparent otherwise) and starts by those with higher chances of being part of an attack.

4 TRAFFIC-AGGREGATE INFERENCE

We now describe the theoretical basis behind *ACC-Turbo*'s inference component. First, we phrase the problem formally and propose a practical solution based on online clustering (§4.1). Second, we introduce the design decisions that make *ACC-Turbo* implementable in existing programmable switches (§4.2). The resulting clustering algorithm can be found in Appendix B.

4.1 Problem definition

Let us define a packet p as a set of features $\mathcal{F}$, where each feature corresponds to a field from the packet header (e.g., *sport, dport, ip.ttl, ip.proto*). For each feature $f \in \mathcal{F}$, packet p has a specific value associated: p_f. We distinguish two types of features: *ordinal* features, for which closer proximity between their values implies stronger similarity (e.g., *ip.src, ip.dst, ip.len, ip.ttl*), and *nominal* features, for which closer values do not necessarily imply similarity (e.g., *sport, dport, ip.proto*).

Let us define an aggregate a as the same set of features $\mathcal{F}$. For each *ordinal* feature f, the aggregate a has a *range* of values associated: $f(a) = [min_f(a), max_f(a)]$. For each *nominal* feature f, aggregate a has a *set* of discrete values associated $f(a) = \{min_f(a), ..., max_f(a)\}$. With this definition, aggregate a represents *all* the packets with feature values included in its ranges and sets.

Objective. At the high-level, given a set of incoming packets, our goal is to infer a set of aggregates that represent *all* the observed packets as *precisely* as possible. We need to limit the number of aggregates to infer by a parameter, $|\mathcal{A}|$. Otherwise, one could list each packet as a separate aggregate and obtain perfect precision.

Definition 4.1 (Aggregate-inference problem). Let $\delta_f(a)$ be the cost of an aggregate a, for feature f, which measures the number of values that it represents. For ordinal features, $\delta_f(a) = max_f(a) - min_f(a)$. For nominal features, $\delta_f(a) = |f(a)|$ (i.e., the number of values in the set). Let $\delta(a) = \prod_{f \in \mathcal{F}} \delta_f(a)$ be the cost of the entire aggregate, as the product of each individual feature cost. For a given feature f, a set of packets $\mathcal{P}$, and a limit on the number of inferred aggregates ($|\mathcal{A}|$), find the set of aggregates $\mathcal{A}^* = a_1, ...a_{k'}(k' \leq |\mathcal{A}|)$ that represent $\mathcal{P}$ and minimize $\delta_f(\mathcal{A}^*) = \sum_{a \in \mathcal{A}^*} \delta_f(a)$.

The presented problem is equivalent to the intent-inference problem in [31], which is NP-hard. As such, we propose a heuristic solution based on online clustering. Our solution aims at approximating the optimal result while enabling per-packet processing. The cost function in Def. 4.1 estimates the number of different packets that a represents, but it is also a measure of similarity of the packets in a. Indeed, packets that can be represented in a narrow aggregate have higher similarity than those which require a broader aggregate. With this intuition, we build an online-clustering algorithm that groups similar packets by minimizing the cost in Def. 4.1.

Definition 4.2 (Online-clustering framework [13]). For a sequence of points p in $\mathcal{P}$, maintain a collection of $|C|$ clusters such that, when each input point p is presented, either it is assigned to one of the current clusters or it starts off a new cluster while two existing clusters are merged into one.

The online-clustering framework is characterized by an endless stream of data, where each data point (i.e., packet) is seen only once:
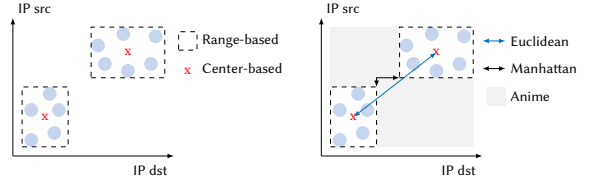

Figure 5: Cluster representations and distances.

it comes in, is processed, and then it goes away, never to return. As such, the algorithm is required to take an irrevocable action after the arrival of each point [25]. Note the difference to the *streaming* case, where the job is finite, and the algorithm can do a second pass through the data to fine-tune the clustering result [9, 21].

4.2 Design decisions

The proposed framework allows three design decisions: (i) the number of clustering possibilities that should be evaluated at each iteration, (ii) the type of information that should be stored about each cluster, and (iii) the distance that should be used to assess the clustering decisions. We make these decisions with the goal of maximizing performance while staying within the resource constraints of existing data planes (to achieve an in-network design).

4.2.1 Clustering search (fast vs. exhaustive). When a new packet arrives, the online-clustering algorithm can either: (i) merge the new packet to its closest cluster or (ii) merge two existing clusters and create a new cluster for the new packet.

[✓] **Fast search**. If the clustering algorithm only supports step (i), we call it *fast*. This approach follows a linear search and requires only $|C|$ distance computations: one for each existing cluster. *ACC-Turbo* relies on this type of search because it can be implemented on programmable data planes.

[X] **Exhaustive search**. If the clustering algorithm supports steps (i) and (ii), we call it *exhaustive*. This approach follows a quadratic search and requires $\binom{|C|}{2}$ additional distance computations. For a given clustering decision, the exhaustive approach outperforms since its search space includes (but is not limited to) the search space of the fast approach. However, the exhaustive approach is not implementable *at line-rate* in existing programmable data planes: it requires accessing multiple times each cluster's information, while registers in today's pipelines can only be accessed once per packet.

4.2.2 Cluster representation (ranges vs. center). A naive way to represent a cluster is as a mere collection of packets. This is, keeping track of *all* the packets and their respective feature values. Such representation is clearly not practical since it does not scale. Therefore, we study two alternative cluster representations.

[X] **Center-based representation**. We can represent clusters by just a single point (e.g., the center of the cluster). The advantage of this representation is that the distance computation is simple, and centers can be easily updated following some pre-defined learning rate [55]. However, we lose a lot of information, such as how big the cluster is or which packets does it contain, which is useful for cluster assessment and traceability, respectively (cf. §10).

[✓] Range-based representation. Following the problem formulation in §4.1, *ACC-Turbo* represents each cluster c with a range of values for each ordinal feature $[min_f(c), max_f(c)]$, and a set of discrete unique values for each nominal feature $\{min_f(c), ..., max_f(c)\}$. Ranges (resp. sets) represent the feature values of packets in a cluster (Fig. 5). This representation preserves information about the cluster sizes, and simplifies interpretability by providing the exact mapping of packets to clusters. Further, ranges and sets are easy to compute and update, which facilitates its implementation on programmable data planes. For instance, the ranges of a new cluster that merges two existing clusters c_i and c_j for feature f are: $[\min(min_f(c_i), min_f(c_j)), \max(max_f(c_i), max_f(c_j))]$. Sets can be implemented as admission lists, using bloom-filters (cf. §6).

4.2.3   The distance function. We first introduce two distance functions for reference, and then derive the distance function that we use in *ACC-Turbo*. For the sake of simplicity, we illustrate the case in which we only consider ordinal features.

[X] Anime distance [31]. The cost function in Def.4.1 can be translated to the online-clustering framework as:

$$\delta_{Anime}(C) = \sum_{c_i \in C} \delta(c_i) = \sum_{c_i \in C} (\prod_{f \in \mathcal{F}} max_f(c_i) - min_f(c_i)) \quad (1)$$

This cost function sums the cost of each individual cluster, which accounts for the number of packets that the cluster represents, and estimates the similarity across packets in the cluster. We now derive a distance function that can be used to assess clustering decisions while trying to minimize this cost. Assuming a range-based cluster representation, we define the distance between clusters c_i and c_j, $\delta(c_i, c_j)$, as the amount of increase in cost produced by merging the two clusters, compared to the total cost of the two clusters if they are not merged: $\delta(c_i, c_j) = \delta(c_i \cup c_j) - (\delta(c_i) + \delta(c_j))$.

The biggest drawback of the Anime distance is the size of the output space. We can measure it as the product of the maximum ranges for each feature. For instance, with the following features and sizes {*ip.len (16b), ip.id (16b), ip.f_offset (13b), ip.ttl (8b), ip.proto (8b), ip.src (32b), ip.dst (32b), sport (16b), dport (16b)*}, the maximum cost is 2^{157}. Data structures in existing data planes can store blocks of maximum 64 bits, which is not enough to represent such cost.

[X] Euclidean distance. Alternatively, we can derive a cost function from the broadly-used Euclidean distance [55]:

$$\delta_{Euclid.}(C) = \sum_{c_i \in C} \delta'(c_i) = \sum_{c_i \in C} \left(\sum_{f \in \mathcal{F}} \sum_{p \in c_i} \|p_f - r_f(c_i)\|^2 \right), \quad (2)$$

where $r_f(c_i)$ is the representative of cluster c_i for feature f. In this case, the output-space size is much smaller. For the same example, the maximum cost can be represented with less than 20 bits. However, computing the Euclidean distance involves square and root operations, which are not straightforward for existing data planes.

[✓] Manhattan distance. In *ACC-Turbo*, we use an alternative distance metric that gathers the benefits of both. Starting from the Anime distance and trying to reduce the size of its output space, we substitute the product of all the individual feature distances by a summation. With this modification, the overall cost function to be minimized can be written as:

$$\delta_{Manh.}(C) = \sum_{c_i \in C} \delta''(c_i) = \sum_{c_i \in C} (\sum_{f \in \mathcal{F}} max_f(c_i) - min_f(c_i)) \quad (3)$$

The resulting distance function to compare a new packet p with one of the clusters, c_i, can be written as: $\delta(p, c_i) = \delta(p \cup c_i) - (\delta(p) + \delta(c_i))$. We know that $\delta(p) = 1$ will have a fixed value, so it will not impact the comparison. First, $\delta(p \cup c_i) = \sum_{f \in \mathcal{F}} \delta_f(p \cup c_i)$, where:

$$\delta_f(p \cup c_i) = \begin{cases} max_f(c_i) - p_f, & \text{if } p_f < min_f(c_i), \\ p_f - min_f(c_i), & \text{if } p_f > max_f(c_i), \\ max_f(c_i) - min_f(c_i), & \text{otherwise.} \end{cases} \quad (4)$$

Second, we can write $\delta(c_i) = \sum_{f \in \mathcal{F}} \delta_f(c_i)$, where $\delta_f(c_i) = [max_f(c_i) - min_f(c_i)]$. Finally, $\delta(p, c_i) \approx \delta(p \cup c_i) - \delta(c_i) = \sum_{f \in \mathcal{F}} [\delta_f(p \cup c_i) - \delta_f(c_i)] = \sum_{f \in \mathcal{F}} \delta_f(p, c_i)$, where:

$$\delta_f(p, c_i) = \begin{cases} min_f(c_i) - p_f, & \text{if } p_f < min_f(c_i), \\ p_f - max_f(c_i), & \text{if } p_f > max_f(c_i), \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

What results from simplifying the Anime distance into a one-dimensional space is the Manhattan distance from the packet to its closest point of the cluster. For a single dimension (i.e., feature), the Manhattan distance and the Anime distance are equivalent. For higher dimensions, the Manhattan distance compresses the output space. As such, it loses information with respect to the original Anime distance. On the other hand, it is easier to compute. Further, it generates outputs in the linear space (as we have sums instead of products), becoming implementable in existing data planes.

5   CONTROLLING AGGREGATES

We now describe how *ACC-Turbo* uses programmable scheduling to mitigate attacks. First, we introduce the programmable-scheduling design space (§5.1) and then, *ACC-Turbo*'s scheduler (§5.2).

5.1   Programmable-scheduling design space

Programming a scheduler consists in defining the order in which it should drain packets from a given buffer. This is done by tagging each packet with a *rank* that indicates the priority with which it should be drained.[3] These ranks are then processed by a (programmable) scheduler that dequeues the packets trying to follow the order specified [24, 45, 56].

Ranking algorithms for pulse-wave attacks. From our definition in §4, a ranking algorithm to mitigate pulse-wave attacks should deprioritize aggregates of high bandwidth and high packet similarity. A number of ranking algorithms can be proposed with this criteria. For example, $rank(p) = throughput(c_i)$, $rank(p) = num.packets(c_i)$, and $rank(p) = throughput(c_i)/size(c_i)$ deprioritize packets by throughput, packet rate, and a combination between throughput and cluster size, respectively. These algorithms can be easily computed with the available data-plane resources. Indeed, aggregate rates can be obtained from packet counters, and packet similarity can be extracted from cluster sizes (cf. §4).

---

[3]Generally, a lower rank indicates a higher priority.

5.2 Scheduling algorithm

To use the proposed ranking algorithms for pulse-wave DDoS defense *today*, we need to make them fit into the limited resources of existing data planes. This is hard for three reasons. First, these resources need to be shared with *ACC-Turbo*'s online-clustering module, which is resource exhaustive (e.g., 12 stages for 4 clusters and 4 features (§6)). Second, existing programmable switches do not support schedulers able to process ranks, forcing us to "build" our own. While it is possible to approximate this scheduling logic using priority queues [24, 56], doing so requires additional resources (e.g., one stage per priority queue [24]). Third, the clustering and scheduling modules must operate sequentially since the ranking algorithm requires the clustering results.

Design. We build a programmable scheduler on top of priority queues and offload the rank computation and the queue mapping to the control plane. Specifically, the control plane periodically (i) polls information about the extracted clusters from the data plane, (ii) assesses clusters' maliciousness and maps them to a priority queue, and (iii) deploys this mapping into the data plane such that *future packets* of each cluster can be scheduled accordingly. This design dedicates all the data-plane resources to the inference process, maximizing its accuracy and preserving line-rate processing.

6 IMPLEMENTATION

We implement *ACC-Turbo* in P4$_{16}$ [11] on Intel Tofino Wedge 100BF-32X [1], with 2484 lines of code. Our prototype uses 12 stages and supports 4 features and 4 clusters. For each incoming packet, *ACC-Turbo* computes its distance to each cluster, selects the closest cluster, and enqueues the packet with the selected cluster's priority.

Cluster selection. For each cluster c, we store the minimum and maximum values of its ordinal-feature's ranges $[min_f(c), max_f(c)]$ using registers. We store its nominal-feature's sets using an admission list, implemented as a bloom filter (cf. §4).

For each arriving packet, we compute its distance to each cluster by computing all the per-cluster per-feature distances, and aggregating them per cluster (cf. Appendix B). For *ordinal* features, we compute the per-cluster per-feature distances by, first, accessing the register containing $min_f(c)$, and checking if the packet feature, p_f, is below the stored value. If that is the case, we set the distance to $d_f(p, c) = min_f(c) - p_f$ within the register's ALU. If not, we access the register containing $max_f(c)$ and set the distance to $d_f(p, c) = p_f - max_f(c)$ if $p_f > max_f(c)$. Since the two registers are accessed sequentially, this takes two stages. However, computation for different cluster-feature pairs can be parallelized. For *nominal* features, we set the distance to 1 if the bloom-filter entry matched by p_f is empty. This takes one stage.

We aggregate the per-cluster per-feature distances into per-cluster distances by progressively summing two distances at each stage using non-stateful ALUs. This requires $log_2|\mathcal{F}|$ stages, being $|\mathcal{F}|$ the number of features. Finally, we find the minimum distance by progressively comparing two distances at each stage also using non-stateful ALUs. This requires $log_2|C|$ stages, being $|C|$ the number of clusters. The distance-computation and distance-aggregation operations can also be parallelized to reduce the number of stages.
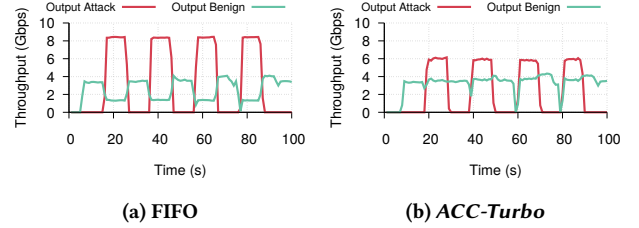

(a) FIFO (b) *ACC-Turbo*

Figure 6: Mitigation of a pulse-wave DDoS attack.

Cluster update. When the minimum distance is not zero, we know that the packet has fallen outside of the selected cluster's coverage. In that case, we need to update the cluster's ranges and sets to accommodate the new packet. We do so by using resubmission.

Queue selection. We use a match-action table, populated by the controller, to map the packet into a priority queue based on the packet's selected cluster. The controller defines the cluster priorities from cluster statistics polled from the data plane (i.e., register entries and packet counters) following the scheduling policy in §5.

Resource requirements. Our implementation (in Tofino 1 [1]) supports a limited number of clusters and features due to Tofino's limited number of stages. Newer programmable switches (e.g., Tofino 2 and 3 [28, 40]) have a higher number of stages, allowing more-performant implementations with more clusters and features.

7 HARDWARE-BASED EVALUATION

We evaluate our hardware implementation of *ACC-Turbo* on Tofino. First, we evaluate *ACC-Turbo*'s performance under a pulse-wave DDoS attack (§7.1). Second, we compare *ACC-Turbo*'s performance to the one of Jaqen [35], a state-of-the-art DDoS defense (§7.2).

7.1 *ACC-Turbo*'s performance

We generate traffic between two servers, connected by a Tofino switch, using interfaces of 100 Gbps (sender→Tofino) and 10 Gbps (Tofino→receiver). As in previous work [35, 58], we replay CAIDA traces as background traffic [12] and add attack traffic on top using MoonGen [22]. The pulse-wave DDoS attack that we generate is composed of four UDP-flood pulses, which have a duration of 10 seconds each and are followed by a 10-second interleave. Each pulse targets a different IP address within a common subnet and a different port. At its peak, the traffic reaches 40.789 Gbps. We configure *ACC-Turbo* to use 4 clusters and to use the last two bytes of the IP destination address, the source port, and the destination port as clustering features. We use a throughput-based ranking algorithm and update the cluster priorities at the controller's maximum speed.

Recovery rate of background traffic. Fig. 6a shows the traffic throughput evolution under no protection. The attack severely impacts the background traffic, with a throughput reduction of ≈ 61%. Note that we are replaying traffic traces and do not see the impact of end-host congestion control. With the effect of congestion control, performance would worsen even further. When we use *ACC-Turbo*, as soon as it infers the attack and deprioritizes its traffic, background traffic fully recovers its original throughput (Fig. 6b).

	Detection	Reaction	Mitigation
ACC-Turbo	Clustering	Always-on	Programmable scheduling
Jaqen [35]	Signature (sketches)	Threshold-based	Rate limit / Drop

Table 2: *ACC-Turbo*'s techniques vs. Jaqen's.

Benign packet drops (%)	FIFO	Jaqen[†]	Jaqen[‡]	*ACC-Turbo*
No Attack	0.00	2.50	3.68	0.00
Single Flow	89.85	2.67	3.95	14.79
Carpet Bombing	89.88	73.19	3.95	19.98
Source Spoofing	89.87	88.16	88.51	14.80

Table 3: Mitigation efficiency under attack variations.

Reaction time. Fig. 6b illustrates the range of possible reaction times in *ACC-Turbo*. While *ACC-Turbo* reacts to the first pulse almost immediately, it takes up to $\approx$ 1s to react to the last two pulses. This reaction time is the time it takes for the control plane to (i) poll the throughput of each cluster, (ii) update their priorities, and (iii) deploy them to the data plane. In our testbed, we leverage a non-optimized Python controller, which takes several milliseconds.

7.2 Comparison to the state-of-the-art DDoS-mitigation technique

We compare *ACC-Turbo* to Jaqen [35], the state-of-the-art DDoS defense. Jaqen uses sketch-based signatures to detect attacks and rate-limiting/dropping to mitigate them (cf. Table 2). With respect to Jaqen, we show that *ACC-Turbo* …

(1) …is more generic (§7.2.1), since it infers attacks agnostically, and not relying on pre-configured signatures.

(2) …achieves faster reaction time (§7.2.2), since it can mitigate attacks without reprogramming the switch.

(3) …does not suffer from the "threshold-based activation" vulnerability (§7.2.3), since it runs continuously on all traffic.

We use the same setup as in the previous section but with the four bytes of the destination IP address as features. We replay CAIDA traces as background traffic at twice their speed (reaching $\approx$ 7 *Gbps*) during 100 seconds and run attacks on top using MoonGen [22]. We generate attacks at maximum capacity, reaching up to $\approx$ 99 *Gbps*.

7.2.1 **Genericity**. We evaluate genericity by analyzing Jaqen's and *ACC-Turbo*'s robustness to attack-traffic variations. Specifically, we consider a UDP-flood attack, which initially consists of a single UDP flow (all the packets share the 5-tuple). We then modify the attack traffic by using: (i) UDP carpet bombing [10, 16, 27] (i.e., the attack targets a /24 destination prefix instead of a single IP); and (ii) UDP source spoofing. We configure Jaqen with a sketch that detects heavy hitters either by monitoring the 5-tuple (Jaqen[†]) or the source IP (Jaqen[‡]). We measure the percentage of benign packets dropped and show the results in Table 3.

We observe that Jaqen is only effective when the correct defense is deployed, but its performance drastically decreases as soon as the attack pattern varies. Instead, *ACC-Turbo* performs well generally, being more robust to attack variations. This is expected: while *ACC-Turbo* infers attacks agnostically, without any initial assumption of the attack characteristics, Jaqen is signature-based and relies on pre-configured defenses for a fixed subset of attacks.

7.2.2 **Reaction time**. We evaluate *ACC-Turbo*'s and Jaqen's reaction time, defined as the time since they see the first attack packet until they start mitigating the attack. For *ACC-Turbo*, the reaction time is the time it takes for the control plane to poll the cluster statistics, update the cluster priorities, and deploy them to the data plane. For Jaqen, it is the time it needs to: detect the attack, compute the right mitigation, orchestrate the network to reroute legitimate traffic, replicate the switch state to the controller, and reprogram the switch with the right mitigation and replicated state.

We measure *ACC-Turbo*'s and Jaqen's reaction times. For Jaqen, we consider two cases: (i) when it needs to deploy a new mitigation (worst case); and (ii) when the right mitigation is already in the switch (best case). Our results show that *ACC-Turbo* reacts at least 11× (resp. 10×) faster than Jaqen in the first (resp. second) case.

***ACC-Turbo*'s reaction time**. We generate a simple UDP-flood attack (all packets sharing the 5-tuple) on top of the CAIDA trace (Fig. 7a), and measure the time it takes for *ACC-Turbo* to react. As depicted in Fig. 7b, *ACC-Turbo* takes $\approx$ 1*s* to react. This is the time required by the (here, unoptimized) control plane to poll the cluster statistics, and deploy the updated priorities to the data plane.

Jaqen's reaction time (defense not deployed). We measure how fast can Jaqen deploy a mitigation which is not already in the switch. To do so, we measure how long it takes for a hardware switch to swap between two (trivial) programs, which simply rewrite the source IP of all packets. We execute the first program for one minute before instructing the switch to swap to the second program, which has been pre-compiled and cached. We send traffic continuously and measure its downtime. We repeat the experiment 10 times.

On average, it takes 11.5 seconds for Jaqen to deploy a new mitigation. This is 11× slower than *ACC-Turbo*'s reaction time. Fig. 7c shows the result for one of the iterations.

Jaqen's reaction time (defense already deployed). If the mitigation module is already loaded, Jaqen's reaction time is the time it takes to detect the attack from the control plane and to deploy a dropping rule to the data plane. We measure this reaction time when Jaqen is configured with the simplest defense (to ensure it is as fast as possible): a sketch that monitors the number of attack packets and a controller that reads it periodically, activating a dropping action when an attack is detected. We read the sketch entries at maximum speed and optimize the threshold value for performance.

As shown in Fig. 7d, Jaqen's reaction time is $\approx$ 10 seconds. This is *still* 10× slower than *ACC-Turbo*. This is the time required by the controller to read the sketch and deploy the drop action and the time it takes for the threshold to be reached–not once, but twice–since Jaqen's only considers attacks when detected in two consecutive time windows.

Jaqen's slow reaction time, of $\approx$ 10 seconds, even in the ideal case, makes it vulnerable to pulse-wave DDoS attacks.

(a) FIFO

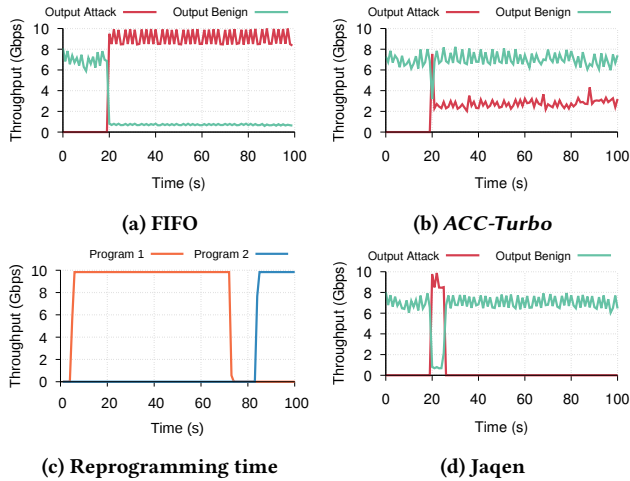**(b) ACC-Turbo**

(c) Reprogramming time

(d) Jaqen

Figure 7: Reaction-time evaluation.

*7.2.3 **Threshold-configuration sensitivity**.* We analyze how sensitive Jaqen is to the threshold-based defense activation vulnerability introduced in §2.2. To that end, we take Jaqen's simplest possible defense: the 5-tuple heavy hitter in §7.2.1 (Jaqen[†]). Such defense relies on two parameters: the *threshold* over which traffic is considered to be an attack and the *periodicity* at which this threshold is checked. We analyze the mitigation efficiency when the two parameters are modified, in the case of a simple UDP-flood attack on top of a CAIDA trace. We measure the percentage of benign traffic dropped and compare the results to the case in which no-defense (i.e., FIFO), and *ACC-Turbo*, are used.

Fig. 8a illustrates Jaqen's high sensitivity to threshold configuration. Slight variations in the threshold value (e.g., from $1M$ packets to $7M$ packets) can result in a drastic increase of benign-packet drops (from $\approx 10\%$ to $\approx 75\%$). This is expected and aligned with our arguments in §2.2. Indeed, the best threshold value depends on the dynamics of both attack and benign traffic, which continuously change over time. As also expected, too-low thresholds can result even worse than not having any defense. This is because they may drop benign traffic even in case of no congestion.

Fig. 8b shows how the periodicity at which the threshold is checked can further impact a certain threshold's performance. For example, a threshold of $10^4 packets$, which outperforms at the controller's maximum periodicity, performs very poorly if the periodicity decreases. At the same time, a bad-performing threshold at maximum periodicity (e.g., $10^7 packets$) can perform well if the controller's periodicity decreases.

ACC-Turbo avoids the threshold-based vulnerability by running continuously on all traffic. Further, *ACC-Turbo* does not perform *binary* assessments on whether a traffic aggregate is malicious or not based on its absolute traffic rate. Instead, it performs *finer-grained* assessments that deprioritize traffic aggregates based on their relative rates (with respect to other aggregates) and their cluster statistics. As a result, even though *ACC-Turbo*'s performance is not as good as Jaqen's when Jaqen is configured optimally, it manages to outperform when Jaqen is not perfectly tuned.
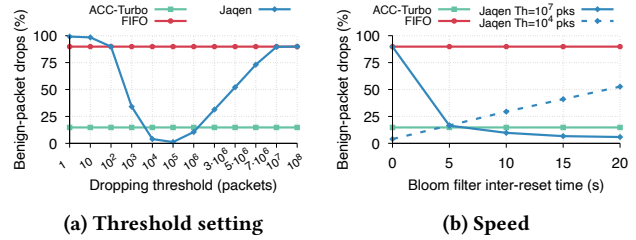

(a) Threshold setting

(b) Speed

Figure 8: Threshold-configuration sensitivity.

8 SIMULATION-BASED EVALUATION

Given the limitations of our Tofino prototype (§6), we extend the evaluation of *ACC-Turbo* by using packet-level simulations in Netbench [3, 30]. We evaluate the performance of *ACC-Turbo* when it schedules a mix of benign traffic and DDoS attacks. We analyze the impact of its different design decisions (§4) and study the performance of more-complete implementations of *ACC-Turbo* (e.g., by using Tofino2 or Tofino3 [28, 40]). We first characterize the clustering strategy (§8.1) and then, the scheduling scheme (§8.2).

Methodology. We simulate an *ACC-Turbo*-enabled switch that processes one day of traffic in which a series of DDoS attacks (or a single morphing attack) is received. We do so by feeding the CICDDoS-2019 trace [44], which contains a sequence of DDoS attacks, into a simulated switch running *ACC-Turbo*. By default, we configure *ACC-Turbo* to support 10 clusters, and to use each byte of the $ip.src$ and $ip.dst$, $sport$, $dport$, $ip.ttl$, and $ip.len$ as features. We adjust the capacity of the output link to various congestion levels.

Overall performance. Even though *ACC-Turbo*'s performance depends on the characteristics of benign and attack traffic, we now illustrate its practicality by evaluating it on a realistic dataset covering a wide range of attack vectors. For all attacks, *ACC-Turbo*'s online-clustering algorithm manages to distinguish packets from attack and benign distributions (achieving cluster's purity of $\approx$ 90%), with as few as 10 clusters. *ACC-Turbo*'s performance is better for well-defined traffic aggregates. For reflection attacks with high packet similarity, *ACC-Turbo* manages to save up to 29% more of benign traffic than FIFO queues, just 5.13% away from an ideal scheduler with full attack knowledge (cf. Fig. 11b, at 50Mbps).

8.1 Characterizing the clustering strategy

We evaluate *ACC-Turbo*'s inference by measuring the *purity* and *recall* of the extracted clusters. These metrics measure the accuracy of the clustering algorithm in mapping packets from different distributions into distinct clusters. We compute purity by (i) labeling each cluster as either majority-benign or majority-malicious, based on the number of clustered packets of each type, (ii) counting the number of packets that match the cluster's label, and (iii) dividing it by the total number of packets [37]. We compute *recall* of benign (resp. malicious) packets as the percentage of benign (resp. malicious) packets mapped into majority-benign (resp. -malicious) clusters. We compute the metrics every one minute and average the result. We only count periods with both attack and benign traffic.
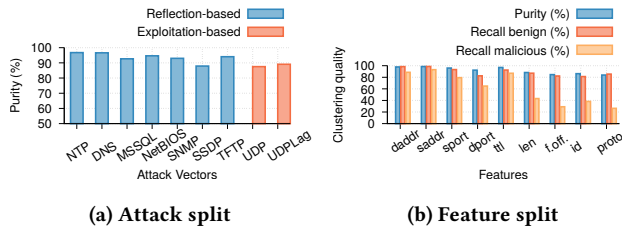
(a) Attack split

(b) Feature split

Figure 9: Performance by attack type and features.
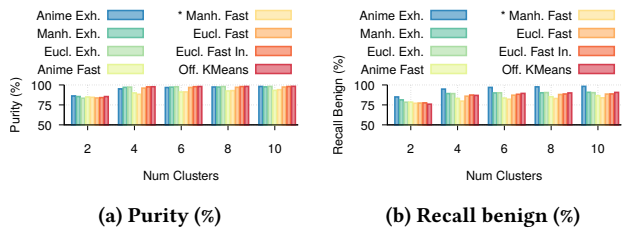

(a) Purity (%)

(b) Recall benign (%)

Figure 10: Performance of clustering strategies.

Feature selection and attack vectors. Fig. 9a shows the clustering performance across attack vectors. In all cases, the achieved purity is above 87%. The clustering performance is strongly related to the variance of the attack features. For example, reflection-based attacks achieve, on average, 5.4% better purity than exploitation-based attacks. Within reflection-based attacks, MSSQL and SSDP, which have higher feature-value variance, perform worst (e.g., MSSQL uses multiple source ports, while NTP or DNS use a single port).

Fig. 9b shows the performance of clustering on individual features. For this particular dataset, IP addresses and source port are good identifiers of malicious traffic. In contrast, fields like IP protocol are less useful as attacks use both UDP and TCP. While the absolute values are tied to the trace characteristics, the split illustrates the different performance levels that *ACC-Turbo* can achieve. While "narrow" attacks achieve good performance if we look at the right features, their performance drastically drops as soon as attacks become diverse or features do not provide a clear signature.

Number of clusters. Fig. 10 shows the performance of different clustering strategies when the number of clusters varies from 2 to 10. First, as expected (cf. §4), a higher number of clusters provides better purity and recall.[4] While selecting the optimal number of clusters is typically a challenge, in *ACC-Turbo* it is imposed by the hardware constraints (§6). Second, increasing the number of clusters is more beneficial for fewer clusters (e.g., the purity in *ACC-Turbo* improves by 4% when we move from 2 clusters to 4, while it only improves 1% when we move from 8 to 10). Since *ACC-Turbo* is designed to run in an environment where the number of clusters is limited, it builds upon this insight and dedicates all data plane resources to maximize the number of clusters (running the non-inference operations in the control plane).

---

[4]In fact, a naive way to achieve perfect purity is to have as many clusters as packets, mapping each packet to its own cluster. From a scheduling perspective, however, this approach would have no value since clusters would not give any information about the maliciousness of the packets contained.

Clustering search: *fast* vs. *exhaustive*. As expected (cf. §4.2.2), we observe that *exhaustive* approaches generally outperform their respective *fast* versions, even though the difference gets smaller as the number of clusters increases. This is especially clear in the Anime and Manhattan approaches, which use range-based representations and have more information to assess clustering decisions (e.g., 93.24% to 98.09% purity increase for Anime with 10 clusters).

Cluster representations and distances. Overall, center-based distances "suffer" less when downgraded from exhaustive to fast (e.g., just 0.79% purity decrease with 10 clusters). There are two reasons for that. First, since they carry little information about clusters, the potential improvement of checking more combinations is limited. Second, range-based approaches are more sensitive to updates since they directly include the new points as they are analyzed. In center-based approaches, we just move the center in the direction of the new data point using some learning rate.

Fast vs. offline and baselines. We compare the performance of *ACC-Turbo* to the one of *offline* k-means with unlimited resources. In all cases, *ACC-Turbo* is very close to the *offline* case (e.g., 4.19% of difference in purity for 10 clusters, cf. Fig. 10). We also study the performance of a hybrid approach, which periodically computes the cluster centers offline and updates them online with the new packets. While the hybrid approach outperforms, the improvement is not significant enough to justify its required increase in complexity.

8.2 Characterizing the scheduling scheme

Finally, we study *ACC-Turbo*'s scheduling performance. First, we evaluate the schedulers in §5, analyzing how often they prioritize benign traffic over malicious traffic. We measure a *score*, defined as the percentage of one-second intervals in the simulation where the average priority given to benign traffic is higher than the one given to malicious traffic. Second, we measure the number of benign packets dropped when the trace is scheduled by a FIFO queue, *ACC-Turbo*, and an ideal scheduler, which prioritizes benign traffic. We use 10 clusters and the 10 most representative features for the trace.

Ranking algorithms. Fig. 11a shows the performance of the scheduling algorithms under the two most complex reflection attacks (cf. Fig.9a). While the absolute values are specific to the dataset, we see how adding the similarity factor to the rank definition (i.e., the cluster sizes) improves performance. This result strongly supports *ACC-Turbo*'s design decision of using a range-based representation.

Bottleneck. Fig. 11b analyzes the number of dropped packets for various bottleneck capacities. First, we see that the *ACC-Turbo* version that is implementable today (i.e., Manhattan distance, fast approach) performs almost on par with the ideal case for smaller bottlenecks. Indeed, at 50 Mbps it saves 29% more benign traffic than FIFO queues, being just 5.13% worse than an ideal PIFO with the ground truth. Second, we compare its performance to the one of Manhattan-exhaustive, and Anime-fast approaches. We see how the loss in performance of today's *ACC-Turbo* comes from the two design decisions required to make it fit into existing devices: using a fast approach and reducing the size of the distance space (§4). With newer programmable switches, some of which are available [28, 40], more-complete versions of *ACC-Turbo* become implementable.
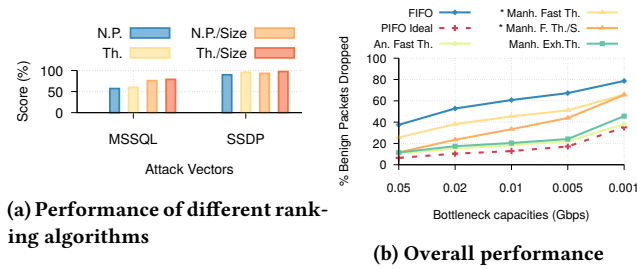
(a) Performance of different ranking algorithms

(b) Overall performance

Figure 11: Impact of scheduling for mitigation.

9 LIMITATIONS

In this section, we discuss techniques by which attackers could try to evade *ACC-Turbo* (§9.1), and mechanisms by which attackers could try to weaponize *ACC-Turbo* to block benign traffic (§9.2).

9.1 Evading *ACC-Turbo*

ACC-Turbo infers attack traffic in an unsupervised fashion, looking for *unexpectedly-high rates* of *traffic aggregates* (i.e., packets that share some similarity). Based on this premise, an attacker willing to evade *ACC-Turbo* has potentially two options: (i) trying to generate low-rate attack traffic and/or (ii) trying to diversify attack packets.

Pulse-wave DDoS attacks are volumetric by nature (i.e., they need *high* traffic rates in order to congest the target link). As such, the first option *per se* is not actually viable. This only leaves the attacker one option: breaking the packet similarity. In the following, we discuss two granularities at which the attacker can try to break packet similarity: at the packet level and at the aggregate level.

Breaking packet similarity at packet level. The goal of this approach is to make it harder for *ACC-Turbo* to correlate attack packets, by adding randomness to one or multiple packet features. Various techniques serve this purpose: e.g., carpet bombing or IP spoofing [10, 16, 27]. These techniques spread the feature values of attack packets all over the space, making *ACC-Turbo* unable to identify any relation among them. While *ACC-Turbo* is robust to a certain degree of randomness (cf. §7), it can not infer attack traffic if all the clustering features are randomized. In the worst scenario, this approach can end up with attack packets mapped to all the clusters and with *ACC-Turbo* not being able to mitigate the attack.

Given that these techniques are well known and widely used, we propose two solutions to prevent them. First, network operators can pick clustering features strategically to avoid these behaviors. For example, carpet bombing can be prevented by clustering longer IP prefixes rather than individual destination IPs [27]. Second, operators can leverage network-monitoring tools to proactively detect and mitigate these patterns, e.g., directly in the data plane or as a part of the cluster's assessment of *ACC-Turbo* in the control plane.

Breaking packet similarity at aggregate level. This approach consists in generating multiple low-rate attack aggregates formed by different traffic types, such that *ACC-Turbo* can not identify any relation between them. In the worst case, this attack can result in *ACC-Turbo* deprioritizing traffic clusters that do not accurately represent attack traffic and not being able to mitigate the attack.

The simplest instantiation of this attack (in number of attack vectors required) is the one composed of $|C|$ spread-out attack vectors, where each vector targets a different cluster. The goal is to attack all the clusters simultaneously, such that *ACC-Turbo* becomes ineffective. While possible in theory, this attack is challenging to execute in practice. Indeed, it requires the attacker to: (i) infer the clustering features and the ranking policy of the victim *ACC-Turbo* instance, (ii) find out which are the attack vectors that maximize their respective distance in the feature space and minimize their probability to be deprioritized by the ranking policy, and (iii) generate these attack vectors, while making sure that they can still reach the victim target. Even though it may be possible to execute this attack for some scenarios, it requires higher complexity than the needed for a conventional DDoS attack. Further, this attack becomes harder to execute, and decreases its effectiveness, linearly to the number of clusters. This is encouraging given the higher number of resources available in newer generations of programmable switches (§7).

9.2 Weaponizing *ACC-Turbo*

We now discuss two attacks which are enabled by *ACC-Turbo*. Here, the goal of the attacker is not to evade *ACC-Turbo*, but to trick it into treating some portion of benign traffic as malicious traffic.

Swapping attack. This attack aims at causing *most* benign (resp. malicious) traffic to be treated as malicious (resp. benign). For that to happen, benign traffic needs to already have a high packet similarity and a high rate. In that case, an attacker can generate lower-rate traffic with randomized packet headers to congest the target link while evading *ACC-Turbo*'s inference. In the worst case, this attack can result in malicious traffic being prioritized over benign traffic. This attack can be perceived as a special case of packet-level evasion (§9.1) and can be defended using the same techniques.

Imitation attack. This attack consists in generating traffic that resembles production traffic, with the goal of triggering *ACC-Turbo* to deprioritize *both* the attack *and* the victim traffic. Same as for previous attack, this attack is challenging to execute in practice. Indeed, attackers need to predict how the victim's traffic will look like at the time of the attack and to be able to accurately replicate it at a high rate (at least for the features supported by *ACC-Turbo*).

The problem of imitation attacks (or attacks composed of indistinguishable flows) has been studied by previous literature in the context of link-flooding attacks [29, 54]. Solutions involve rate-change tests or historical-pattern analysis to shed light on whether a certain aggregate is a legitimate service or a spoofed replica [29].

10 DISCUSSION

Are today's DDoS still "aggregates"? Yes. Even though DDoS attacks can *theoretically* be arbitrarily complex, most DDoS attacks today are still formed by well-defined traffic subsets [27, 32], thus being characterizable as "aggregates" [36]. For instance, the Mirai attack, which is probably the most famous botnet-based attack to date, included several flooding attacks such as UDP flood, SYN flood, ACK flood, or HTTP flood [8]. Each such attack generated highly-similar packets [53]: for instance, all SYN-flood packets had the same size, protocol, flags, and shared the same source IP subnets.

Amplification attacks, such as the NTP-based attack on Cloud-flare (2014), DNS-based attack on Google (2017), or the Memcached-based attack against GitHub (2018), can be generally characterized by unusually-large packets sourcing from a specific port, using the same protocol, and originating from a common subset of IPs [6, 18].

Finally, alleged link-flooding attacks [41, 48] (cf. [29]) were also composed by (i) "DNS responses of 3000 bytes and TCP reflections targeting specific addresses of the victim IXP" [41], and (ii) "ampli-fication vectors such as NTP, UDP, TCP, and ICMP-floods" [48].

Can *ACC-Turbo* cover other attack types? While *ACC-Turbo* has been designed to cover pulse-wave DDoS attacks, it can also mitigate conventional DDoS attacks as long as they are volumetric and composed of clear traffic aggregates (as we show in §8).

On the other hand, *ACC-Turbo* does *not* cover application-layer attacks nor low-bandwidth attacks. Since these attacks are not volumetric, they do not generate congestion in the network and remain unnoticed by *ACC-Turbo*. However, as we discuss in §9, *ACC-Turbo* can be deployed together with complementary defenses, which may protect against attacks that *ACC-Turbo* can not cover.

What is the impact of leaving *ACC-Turbo* always-on? Under sporadic congestion, the possible packet reorderings produced by *ACC-Turbo* could be happening without our solution as well, e.g., if the deprioritized packets suffered from longer delays on the network. In case of sustained congestion (not an attack), *ACC-Turbo* will deprioritize groups of packets with higher rates (heavy hitters) and give more priority to less aggressive groups of packets. The impact could be similar to a fair-queuing scheme, with the difference that the definition of a flow is inferred dynamically.

What about reordering? Assuming that the features used are common across all packets of a given flow, all these packets will be mapped into the same cluster. As such, reordering can only happen when the priority given to the flow's cluster increases over time while there are still packets of the same flow in the queue with the old (lower) priority. Since we update priorities in time windows of milliseconds to seconds, potential reordering would only impact large flows with already-high flow completion times.

What about interpretability? *ACC-Turbo*'s range-based cluster-ing allows operators to know which packets are being mapped into each cluster, as well as the exact mapping of these clusters to the priority queues. Contrary to a black-box approach, an operator can access the complete information of every action performed in real-time. An operator could further modify the table entries in *ACC-Turbo* to reduce the number of priority queues to be used to drop specific parts of traffic or treat some known-benign traffic preferentially (e.g., by mapping it to a dedicated priority queue).

What future research do we envision? Given the limitations of *ACC-Turbo*, we expect future work to tackle two main research di-rections. First, the design of enhanced inference techniques, which can identify attack patterns with higher accuracy and robustness. Second, the combination of signature-based defenses, which out-perform in detecting known attacks (§8), with generic defenses like *ACC-Turbo*, which offer the potential to infer new attacks. The com-bination of both approaches can also result beneficial in preventing adversarial attacks to unsupervised DDoS defenses (cf. §9).

11 RELATED WORK

Given the novelty of pulse-wave DDoS attacks, only a few works have studied their mitigation [15]. Thus, in this section, we cover in-network defenses for conventional DDoS attacks, which are most related to our work. Another group of attacks, called *shrew* or *pulsing attacks*, also leverage traffic pulses to disrupt the victim's connectivity [14]. However, they differ from pulse-wave DDoS attacks because they are low-rate and target TCP vulnerabilities.

Conventional-DDoS defenses. Poseidon [58], Jaqen [35], and Ripple [54] use programmable switches to mitigate DDoS attacks. Poseidon proposes a system-level solution to orchestrate traffic to pre-defined defenses in dedicated servers. Jaqen and Ripple are both signature-based and rely on a network-wide view of attack signals to deploy pre-configured defenses. *ACC-Turbo* runs autonomously in a switch (not requiring network orchestration nor switch re-programming), achieves fast reaction times, and covers unknown attacks. Bohatei [23] uses network function virtualization to adapt the placement, scale, and the type of pre-defined defenses depend-ing on the detected attacks. SPIFFY [29] detects link-flooding attacks by actively modifying the available bandwidth and analyzing the traffic reaction. Kitsune [39] trains an ensemble of autoencoders to identify anomalies in network traffic. Euclid [20] uses statistical analysis to detect attacks from the network. While these solutions achieve high accuracy, they execute drastic mitigation policies such as filtering. *ACC-Turbo* uses programmable scheduling to mitigate attacks directly from the network with reduced collateral damage.

Scheduling-based DDoS defenses. DDoS-Shield [43] uses sched-uling to mitigate DDoS attacks. While *ACC-Turbo* faces pulse-wave DDoS attacks from the network, DDoS-Shield targets application-layer attacks from the end-host. In [33], a DDoS defense is pro-posed, using two priority queues and assigning suspected flows to the low-priority queue. Suspected flows are identified based on the difference in harmonic means of the arrival rate of incoming pack-ets. *ACC-Turbo* uses more priority queues to enable finer-grained decisions and accommodates more complex DDoS attacks.

12 CONCLUSIONS

In this paper, we emphasized the need for in-network pulse-wave DDoS defenses, and showcased the potential of building them by combining unsupervised inference and programmable scheduling.

We presented *ACC-Turbo*, the *first* aggregate-based congestion control mechanism that mitigates pulse-wave DDoS attacks by running at line rate on commodity hardware. *ACC-Turbo* leverages two key insights: online clustering—to effectively identify (possibly unknown) attack pulses, and programmable scheduling—to safely deprioritize traffic according to its maliciousness. We implemented *ACC-Turbo* in P4 and deployed it on programmable hardware. *ACC-Turbo* can mitigate pulse-wave DDoS attacks in almost real-time.

Ethics statement. This work does not raise any ethical issues.

ACKNOWLEDGMENTS

REFERENCES

[1] 2017. Intel Tofino. http://barefootnetworks.com/products/brief-tofino/. (2017).
[2] 2017. Pulse-Wave DDoS Attacks Mark a New Tactic in Q2. https://www.infosecurity-magazine.com/news/pulsewave-ddos-attacks-mark-q2/. (2017).
[3] 2018. NetBench. http://github.com/ndal-eth/netbench. (2018).
[4] 2019. Hidden Threat of Pulse Wave DDoS attacks. https://ddos-guard.net/en/info/blog-detail/hidden-threat-of-pulse-wave-ddos-attacks. (2019).
[5] Link 11. 2020. The Evolution of DDoS Reflection Amplification Vectors: a Chronology. https://www.link11.com/en/blog/threat-landscape/ddos-reflection-amplification-vectors-chronology/. (2020).
[6] Akamai. 2018. Memcached-fueled 1.3 Tbps Attacks. https://blogs.akamai.com/2018/03/memcached-fueled-13-tbps-attacks.html. (2018).
[7] Akamai. 2020. Prolexic Routed: Protect Your Entire Application Infrastructure Against Large, Complex DDoS Attacks. https://www.akamai.com/uk/en/multimedia/documents/product-brief/prolexic-routed-product-brief.pdf. (2020).
[8] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J. Alex Halderman, Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever, Zane Ma, Joshua Mason, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, and Yi Zhou. 2017. Understanding the Mirai Botnet. In *USENIX Security*. Vancouver, Canada.
[9] Albert Bifet, Ricard Gavaldà, Geoff Holmes, and Bernhard Pfahringer. 2018. Machine Learning for Data Streams with Practical Examples in MOA. In *MIT Press*.
[10] Steinthor Bjarnason. 2018. DDoS Defences in the Terabit Era: Attack Trends, Carpet Bombing. https://blog.apnic.net/2018/12/04/ddos-defences-in-the-terabit-era-attack-trends-carpet-bombing/. (2018).
[11] Pat Bosshart, Dan Daly, Glen Gibb, Martin Izzard, Nick McKeown, Jennifer Rexford, Cole Schlesinger, Dan Talayco, Amin Vahdat, et al. 2014. P4: Programming Protocol-Independent Packet Processors. In *ACM SIGCOMM CCR*.
[12] CAIDA. 2018. Anonymized Traces from CAIDA Equinix NYC Internet Data Collection Monitor-B, 2018. https://www.caida.org/data/passive/trace_stats/nyc-B/2018/. (2018).
[13] Moses Charikar, Chandra Chekuri, Tomás Feder, and Rajeev Motwani. 2004. Incremental Clustering and Dynamic Information Retrieval. In *SIAM JoC*.
[14] Chen, Yu and Hwang, Kai and Kwok, Yu-Kwong. 2005. Collaborative Defense Against Periodic Shrew DDoS Attacks in Frequency Domain. In *ACM TISSEC*.
[15] Ilya V. Chugunkov, Leonid O. Fedorov, Bela Sh. Achmiz, and Zarina R. Sayfullina. 2018. Development of the Algorithm for Protection Against DDoS-attacks of Type Pulse Wave. In *IEEE EIConRus*.
[16] Catalin Cimpanu. 2019. 'Carpet-bombing' DDoS Attack Takes Down South African ISP for an Entire Day. https://www.zdnet.com/article/carpet-bombing-ddos-attack-takes-down-south-african-isp/. (2019).
[17] Catalin Cimpanu. 2020. FBI Warns of New DDoS Attack Vectors: CoAP, WS-DD, ARMS, and Jenkins. https://www.zdnet.com/article/fbi-warns-new-ddos-attack-vectors-coap-ws-dd-arms-and-jenkins/. (2020).
[18] Cloudflare. 2014. Technical Details Behind a 400Gbps NTP Amplification DDoS Attack. https://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack/. (2014).
[19] Cloudflare. 2017. Cloudflare Advanced DDoS Protection. https://www.cloudflare.com/media/pdf/cloudflare-whitepaper-ddos.pdf. (2017).
[20] Alexandre da Silveira Ilha, Ângelo Cardoso Lapolli, Jonatas Adilson Marques, and Luciano Paschoal Gaspary. 2020. Euclid: A Fully In-Network, P4-based Approach for Real-Time DDoS Attack Detection and Mitigation. In *IEEE TNSM*.
[21] Sanjoy Dasgupta. 2008. Notes for CSE 291: Topics in Unsupervised Learning, Lecture 6: Online and Streaming Algorithms for Clustering. https://cseweb.ucsd.edu/~dasgupta/291-unsup/lec6.pdf. (2008).
[22] Paul Emmerich, Sebastian Gallenmüller, Daniel Raumer, Florian Wohlfart, and Georg Carle. 2015. Moongen: A Scriptable High-Speed Packet Generator. In *ACM IMC*. Tokyo, Japan.
[23] Seyed K Fayaz, Yoshiaki Tobioka, Vyas Sekar, and Michael Bailey. 2015. Bohatei: Flexible and Elastic DDoS Defense. In *USENIX Security*. Washington, DC, USA.
[24] Albert Gran Alcoz, Alexander Dietmüller, and Laurent Vanbever. 2020. SP-PIFO: Approximating Push-In First-Out Behaviors using Strict-Priority Queues. In *USENIX NSDI*. Santa Clara, CA.
[25] Sudipto Guha, Adam Meyerson, Nina Mishra, Rajeev Motwani, and Liadan O'Callaghan. 2003. Clustering Data Streams: Theory and Practice. In *IEEE TKDE*.
[26] Yuhei Hayashi, Chen Meiling, and Li Su. [n. d.]. *Use Cases for DDoS Open Threat Signaling (DOTS) Telemetry*. Internet-Draft draft-ietf-dots-telemetry-use-cases-07. Internet Engineering Task Force. https://datatracker.ietf.org/doc/html/draft-ietf-dots-telemetry-use-cases-07
[27] Tiago Heinrich, Rafael R Obelheiro, and C. Alberto Maziero. 2021. New Kids on the DRDoS Block: Characterizing Multiprotocol and Carpet Bombing Attacks.
[28] Intel. [n. d.]. Tofino 3 Intelligent Fabric Processor. https://www.intel.com/content/www/us/en/products/network-io/programmable-ethernet-switch/tofino-3-brief.html. ([n. d.]).
[29] Min Suk Kang, Virgil D Gligor, Vyas Sekar, et al. 2016. SPIFFY: Inducing Cost-Detectability Tradeoffs for Persistent Link-Flooding Attacks. In *NDSS Symposium*. San Diego, CA, USA.

[30] Simon Kassing, Asaf Valadarsky, Gal Shahaf, Michael Schapira, and Ankit Singla. 2017. Beyond Fat-trees Without Antennae, Mirrors, and Disco-balls. In *ACM SIGCOMM*. Los Angeles, CA, USA.
[31] Ali Kheradmand. 2020. Automatic Inference of High-Level Network Intents by Mining Forwarding Patterns. In *ACM SOSR*. San Jose, CA, USA.
[32] Oleg Kupreev, Ekaterina Badovskaya, and Alexander Gutnikov. 2020. Kaspersky DDoS reports. DDoS attacks in Q1 2020. https://securelist.com/ddos-attacks-in-q1-2020/96837/. (2020).
[33] Chu-Hsing Lin, Jung-Chun Liu, Hsun-Chi Huang, and Tsung-Che Yang. 2008. Using Adaptive Bandwidth Allocation Approach to Defend DDoS Attacks. In *IEEE MUE*. Busan, Korea.
[34] Thomas Lintemuth, Patrick Hevesi, and Sushil Aryal. 2020. Solution Comparison for DDoS Cloud Scrubbing Centers. In *Gartner*.
[35] Zaoxing Liu, Georgios Nikolaidis, Jeongkeun Lee, Changhoon Kim, Xin Jin, Vladimir Braverman, Minlan Yu, and Vyas Sekar. 2021. Jaqen: A High-Performance Switch-Native Approach for Detecting and Mitigating Volumetric DDoS Attacks with Programmable Switches. In *USENIX Security*. Virtual.
[36] Mahajan, Ratul and Bellovin, Steven M. and Floyd, Sally and Ioannidis, John and Paxson, Vern and Shenker, Scott. 2002. Controlling High Bandwidth Aggregates in the Network. In *ACM SIGCOMM CCR*. New York, NY, USA.
[37] Christopher D. Manning, Hinrich Schütze, and Prabhakar Raghavan. 2008. Introduction to Information Retrieval (Chapter 16.3). In *Cambridge University Press*.
[38] Damian Menscher. 2019. Practical Solutions for Amplification Attacks. https://pc.nanog.org/static/published/meetings/NANOG76/daily/day_2.html. (2019).
[39] Yisroel Mirsky, Tomer Doitshman, Yuval Elovici, and Asaf Shabtai. 2018. Kitsune: An Ensemble of Autoencoders for Online Network Intrusion Detection. In *NDSS Symposium*. San Diego, CA, USA.
[40] Edgecore Networks. [n. d.]. 400 GbE Data Center Spine Switch Bare-Metal Hardware. https://stordirect.com/wp-content/uploads/woocommerce-products-data/product_documents/Edgecore_AS9516-32D_datasheet.pdf. ([n. d.]).
[41] Matthew Prince. 2013. The DDoS That Almost Broke the Internet. https://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet/. (2013).
[42] Radware. 2021. DefensePro: Advanced DDoS Defense and Attack Mitigation. https://www.radware.com/products/defensepro/. (2021).
[43] Supranamaya Ranjan, Ram Swaminathan, Mustafa Uysal, and Edward W. Knightly. 2006. DDoS-Resilient Scheduling to Counter Application Layer Attacks Under Imperfect Detection. In *IEEE INFOCOM*. Barcelona, Spain.
[44] Iman Sharafaldin, Arash Habibi Lashkari, Saqib Hakak, and Ali A Ghorbani. 2019. Developing Realistic Distributed Denial of Service (DDoS) Attack Dataset and Taxonomy. In *ICCST*.
[45] Anirudh Sivaraman, Suvinay Subramanian, Mohammad Alizadeh, Sharad Chole, Shang-Tse Chuang, Anurag Agrawal, Hari Balakrishnan, Tom Edsall, Sachin Katti, and Nick McKeown. 2016. Programmable Packet Scheduling at Line Rate. In *ACM SIGCOMM*. Florianopolis, Brazil.
[46] Jared M. Smith and Max Schuchard. 2018. Routing Around Congestion: Defeating DDoS Attacks and Adverse Network Conditions via Reactive BGP Routing. In *IEEE S&P*. San Francisco, CA, USA.
[47] Robin Sommer and Vern Paxson. 2010. Outside the Closed World: On Using Machine Learning for Network Intrusion Detection. In *IEEE S&P*. Oakland, CA.
[48] Proton Team. 2015. Guide to DDoS Protection. https://protonmail.com/blog/ddos-protection-guide/. (2015).
[49] Proton Team. 2015. Message Regarding the ProtonMail DDoS Attacks. https://protonmail.com/blog/protonmail-ddos-attacks/. (2015).
[50] Alethea Toh. 2022. Azure DDoS Protection—2021 Q3 and Q4 DDoS attack trends. https://azure.microsoft.com/en-gb/blog/azure-ddos-protection-2021-q3-and-q4-ddos-attack-trends/. (2022).
[51] Muoi Tran, Min Suk Kang, Hsu-Chun Hsiao, Wei-Hsuan Chiang, Shu-Po Tung, and Yu-Su Wang. 2019. On the Feasibility of Rerouting-Based DDoS Defenses. In *IEEE S&P*. San Diego, CA, USA.
[52] Jai Vijayan. 2017. 'Pulse Wave' DDoS Attacks Emerge As New Threat. https://www.darkreading.com/attacks-breaches/pulse-wave-ddos-attacks-emerge-as-new-threat-/d/d-id/1329657. (2017).
[53] Ron Winward. 2017. Mirai Inside of an IoT Botnet. https://www.nanog.org/news-stories/nanog-tv/top-talks/mirai-inside-iot-botnet/. (2017).
[54] Jiarong Xing, Wenqing Wu, and Ang Chen. 2021. Ripple: A Programmable, Decentralized Link-Flooding Defense Against Adaptive Adversaries. In *USENIX Security*.
[55] Rui Xu and Don Wunsch. 2008. Clustering (Chapter 5.2.1). In *John Wiley & Sons*.
[56] Zhuolong Yu, Chuheng Hu, Jingfeng Wu, Xiao Sun, Vladimir Braverman, Mosharaf Chowdhury, Zhenhua Liu, and Xin Jin. 2021. Programmable Packet Scheduling with a Single Queue. In *ACM SIGCOMM*. Virtual Event, USA.
[57] Igal Zeifman. 2017. Attackers Use DDoS Pulses to Pin Down Multiple Targets. https://www.imperva.com/blog/pulse-wave-ddos-pins-down-multiple-targets/. (2017).
[58] Menghao Zhang, Guanyu Li, Shicheng Wang, Chang Liu, Ang Chen, Hongxin Hu, Guofei Gu, Qi Li, Mingwei Xu, and Jianping Wu. 2020. Poseidon: Mitigating Volumetric DDoS Attacks with Programmable Switches. In *NDSS Symposium*. San Diego, CA, USA.

APPENDIX

Appendices are supporting material that has not been peer-reviewed.

A   ACC PARAMETERS

Name	Definition	Value
K	Sustained-congestion period	$2s$
p_{high}	Sustained-congestion droprate	0.1
p_{target}	Target droprate	0.05
k	Exponential-moving-average interval for rate estimation	$0.1s$
Sessions	Maximum number of allowed rate-limiting sessions	5
Release Time	Minimum time required for an aggregate to be released after rate-limiting starts	$10s$
Free Time	Minimum time required for an aggregate to be released after it is detected to "behave"	$20s$
Cyc. Time	Time to revisit the aggregate	$5s$
Init. Time	Time to revisit the aggregate in the initial phase	$0.5s$

Table 4: List of the ACC parameters used in the simulations, excluding RED and rate-limiting configurations.

B   *ACC-TURBO* CLUSTERING ALGORITHM

Algorithm 1 ACC-Turbo Clustering Algorithm

Require: p: New packet, min, max: Initial ranges
1: **procedure** CLUSTERING
2:    **for all** p: incoming packet **do**
3:       **for all** $c_i \in C$ **do**
4:          $d(p, c_i) \leftarrow$ COMPUTEDISTANCE(p, c_i)
5:       $c_{selected} \leftarrow c_0$ ▷ Initialize selected cluster
6:       $d_{min} \leftarrow d(p, c_0)$ ▷ Initialize min. distance
7:       **for all** $c_i \in C, \ i \neq 0$ **do**
8:          **if** $d(p, c_i) < d_{min}$ **then**
9:             $d_{min} \leftarrow d(p, c_i)$
10:            $c_{selected} \leftarrow c_i$ ▷ Select cluster
11:      **if** $d(p, c_{selected}) > 0$ **then**
12:         $min, max \leftarrow$ UPDATECLUSTER$(p, c_{selected})$

13:
14: **function** COMPUTEDISTANCE(p, c_i)
15:    $d(p, c_i) \leftarrow 0$ ▷ Initialize distance
16:    **for** $f \in \mathcal{F}$ **do**   ▷ Iterate over all features
17:       $d_f(p, c_i) \leftarrow 0$
18:       **if** f is ordinal **then**
19:          **if** $p_f < min_f(c_i)$ **then**
20:             $d_f(p, c_i) \leftarrow min_f(c_i) - p_f$
21:          **if** $p_f > max_f(c_i)$ **then**
22:             $d_f(p, c_i) \leftarrow p_f - max_f(c_i)$
23:          $d(p, c_i) \mathrel{+}= d_f(p, c_i)$ ▷ Aggregate distances
24:       **else**
25:          **if** $p_f \notin f(c_i)$ **then**
26:             $d_f(p, c_i) \leftarrow 1$
27:    **return** $d(p, c_i)$

28:
29: **function** UPDATECLUSTER$(p, c_{selected})$
30:    **for** $f \in \mathcal{F}$ **do**   ▷ Iterate over all features
31:       **if** f is ordinal **then**   ▷ Update ranges
32:          **if** $p_f < min_f(c_{selected})$ **then**
33:             $min_f(c_{selected}) \leftarrow p_f$
34:          **if** $p_f > max_f(c_{selected})$ **then**
35:             $max_f(c_{selected}) \leftarrow p_f$
36:       **else**   ▷ Update feature-value set
37:          **if** $p_f \notin f(c_{selected})$ **then**
38:             $f(c_{selected}) = f(c_{selected}) \cup p_f$
39:    **return** min, max