A Game Theoretical Analysis of Distributed Denial-of-Service Defense Incentive

Mingwei Zhang[1], Jun Li[1(✉)], Jiabin Wu[1], and Peter Reiher[2]

[1] University of Oregon, Eugene, USA
{mingwei,lijun,jwu5}@uoregon.edu
[2] University California, Los Angeles, USA
reiher@cs.ucla.edu

Abstract. Distributed denial-of-service (DDoS) attacks are becoming more frequent and powerful. Traditional *edge defense* solutions can no longer keep up, and *in-network defense* solutions are needed that involve multiple Internet Service Providers (ISPs) to collaboratively defend against the attacks. While collaborative defense solutions are technically more effective at stopping large-scale attacks, the incentives for ISPs to deploy these solutions remain unexplored. In this study, we develop a game theoretic model to capture the economic benefits and costs of deployment for ISPs competing for customers. Through large-scale simulations at the Internet level, we find that the majority of ISPs on the Internet have an economic incentive to participate in DDoS defense, driven by competition; and that the severity of DDoS attacks and the level of competition affect an ISP's charge for filtering DDoS traffic for its customers.

Keywords: DDoS · DDoS defense · in-network DDoS defense · DDoS defense incentive

1 Introduction

Distributed denial-of-service (DDoS) attacks have plagued the Internet for more than two decades. DDoS attacks use a large number of attack sources (e.g. compromised computers) to flood victims' networks with unwanted traffic in order to exhaust victims' network, computation, and other types of resources. DDoS attacks of more than 1 terabit per second have become common [1,19], which can pose a severe threat to all online services.

Traditionally, DDoS defense has been thought of as *edge defense*, where either the DDoS victim itself or an entrusted third party conducts the defense at the edge of the Internet. There are two problems that make edge defense less effective against current and future DDoS attacks. First, the cost for a single entity to handle DDoS attack traffic at the terabit-per-second level is usually very high. To make matters worse, today's attackers can more easily tap into the increasingly popular but less secure Internet of Things devices to launch high-volume

H. Duan et al. (Eds.): SecureComm 2023, LNICST 568, pp. 3–24, 2025.
https://doi.org/10.1007/978-3-031-64954-7_1

DDoS attacks. Second, even with sufficient investment in handling incoming DDoS traffic at the edge, in many cases the defense is already late due to traffic congestion that occurs before reaching the edge. For example, research in [18] showed that attackers can create traffic congestion around the victim without directly launching an attack at the victim.

So people turn to *in-network DDoS defense*, which, as the name implies, places defense efforts along the paths of DDoS traffic, blocking DDoS traffic from inside the Internet before it reaches the victim. It has the following advantages over edge defense. First, in-network defense allows the defense load to be shared among the defense parties, reducing the defense effort required at each party while potentially achieving a higher overall defense capacity. Second, the filtering of DDoS traffic can be done earlier, reducing the traffic load on the victim while mitigating the traffic congestion on the links before reaching the victim. Overall, in the current Internet environment, in-network DDoS defense appears more suitable and attractive than edge defense.

However, it is still unknown whether in-network defenses are economically beneficial to the defenders in the network. In edge defense cases, the defender is either the victim or directly serves the victim, so the incentive to defend is clear. However, for anyone involved in in-network defense, the incentive to participate is unclear. Filtering DDoS traffic to its customers would reduce an ISP's profit due to the reduction of the amount of traffic it forwards. For example, if an Internet service provider (ISP) is the sole provider for its customers, as long as its revenues exceed its costs, it is always in the ISP's best interest to forward as much traffic as possible.

On the other hand, as the Internet infrastructure grows, no single ISP can monopolize the entire service market. Instead, ISPs compete with each other. Customers who pay for Internet access expect uninterrupted service and will naturally choose the ISP that does the most to stop or mitigate incoming DDoS attacks. By investing more in DDoS defense, an ISP has a better chance of being selected by customers among its competitors, and the right to carry traffic for customers generates revenue. Nonetheless, providing DDoS defense service is not without its costs. First, an ISP would lose revenue by dropping the DDoS traffic that would have been delivered to customers. Second, deploying new defense solutions is expensive, both in terms of equipment and maintenance costs. In a nutshell, while the benefits of winning the competition can incentivize ISPs to invest in DDoS defense, it is also possible that an ISP may choose not to compete due to the potential loss of revenue and increase in workload and costs.

In this study, we propose a game theoretical model to analyze the profit maximization decisions of ISPs who are competing for customers by investing in DDoS defense. We design and implement a simulation framework to simulate the decision process of DDoS defense deployment for all ISPs on the Internet. The simulation results show that we can find an equilibrium, and the majority of ISPs would decide to participate in DDoS defense to maximize their profits. Depending on the severity of DDoS attacks and the level of competition, ISPs can charge different rates to achieve their optimal profits.

The rest of the paper is organized as follows. Section 2 summarizes the related work of this study; Sect. 3 introduces a game theoretical model of competition among ISPs for customers; Sect. 4 describes the simulation setup and algorithms for this study; Sect. 5 discusses the simulation results; and Sect. 6 concludes the paper.

2 Related Work

There is a large body of work that addresses cybersecurity problems from both incentive and game theory perspectives [20,22,26,28]. Some studies look at the interactions between attackers and defenders, while other work explores interdependencies and cooperation among defenders.

Numerous research projects have explored the use of game theory to study interactions between attackers and defenders in the context of cybersecurity. Bedi et al. in [3] proposed a model to study optimal firewall settings for DDoS defense against attackers. Bohawek et al. in [6] introduced game theoretic stochastic routing (GTSR) to minimize the impact of link and router failures against intelligent attackers. Shiva et al. in [32] proposed a holistic architecture that incorporates attacker behavior and decides actions for the defenders, but the work lacks concrete evaluation. Wu et al. [34] developed a game-theoretic model to study the most effective firewall settings to block DoS/DDoS traffic. All of the above work attempts to model and design systems to mitigate attacks (which can be intelligent and dynamic) from a single, centrally controlled entity. However, current cyber attacks on the Internet, especially DDoS attacks, can no longer be easily mitigated by single AS.

There are studies that investigate the potential cooperation among defenders against cyber attacks. Grossklags et al. in [15] analyzed how influences among heterogeneous entities could reach different security outcomes under five different economic environments. Later, they studied how the interdependent defenders could switch between public good (protection) and private good (insurance) given the choices [16]. Similarly, Miura-ko et al. [25] model the impact of security investments among interdependent organizations using an influence network. However, the lack of quantitative evaluation and conclusion makes it less applicable to real-world problems.

Some work studies cybersecurity with a focus on defense incentives. Early work by Huang et al. [17] analyzed the broken incentive chain that stops ISPs from participating in DDoS defense. They argue that the subscription-based pricing model among ISPs at the time, which often leads to overprovisioning and ignores actual traffic volume patterns, discourages ISPs from participating in defense by doing extra work. They suggest that a traffic-usage-based pricing model would incentivize ISPs to help filter out unwanted traffic. Unfortunately, the authors did not consider the potential revenue that the ISPs could have made by *not participating in DDoS defense*; therefore, as the Internet gradually moves to a usage-based pricing model, the apparent lack of incentive persists. Gill et al. [13] argue that efforts to deploy more secure inter-domain routing protocol would

attract more inter-domain traffic, and thus generate more revenue, consistent with the results of the earlier discussion in [29]. Unlike DDoS defense, securing inter-domain routing has no disincentive, i.e., not defending does not introduce additional revenue. However, as the authors suggest, both strong early adopters and a simplified protocol are needed for global deployment. Shen et al. [31] use game theory to investigate the deployment incentive in their previous work They model the deployment problem as a social dilemma, and suggest that ISPs can be incentivized by combining the benefits of achieving the public good and the potential punishment for untrustworthy behavior. However, we believe that public good is not a strong motivation for private firms, and that trust assessment and behavioral punishment require enforcement by global central authorities, which is not a realistic assumption under the context of today's Internet.

3 Game of DDoS Defense Investment

In this section, we introduce a game-theoretic model to capture the strategic interactions between parties on the Internet that are involved in DDoS defense and analyze their profit and cost in DDoS defense.

3.1 Network Modeling

Internet Topology. The Internet is composed of interconnected **Autonomous Systems (ASes)**, where each AS decides autonomously how its networks are connected to each other and to other ASes. Every ISP is one or more ASes. We consider the Internet as a weighted directed graph $G = \{V, E\}$. V represents the set of all nodes (vertices) in the graph, where every node $v_i \in V$ represents an AS on the Internet. E represents the set of all edges in the graph, where every edge $e_{i,j} \in E$ represents an inter-AS directed link between two neighboring ASes v_i and v_j.

ASes form business relationships to establish links (both physical and topological) to connect themselves to other entities and further to the rest of the Internet. As summarized by Gao [12], there are three common types of business relationships between ASes: customer-to-provider, peer-to-peer, and sibling-to-sibling. In a customer-to-provider relationship, the provider AS carries traffic to and from the customer AS, and the customer AS pays its provider AS for the traffic the provider carries for it. In a peer-to-peer or sibling-to-sibling relationship, two ASes forward traffic for each other, usually for free. In this study, we mainly consider the ASes with customer-to-provider relationship.

For any IP prefix (i.e., a block of IP addresses with a common prefix) to be reachable by end hosts connected to the Internet, the AS that owns the IP prefix must advertise the IP prefix to the Internet using the Boarder Gateway Protocol (BGP). BGP is used to propagate the prefix reachability information across the Internet between each pair of neighboring ASes. Upon receiving an advertised path to a prefix, an AS decides whether to use the path for reaching the prefix

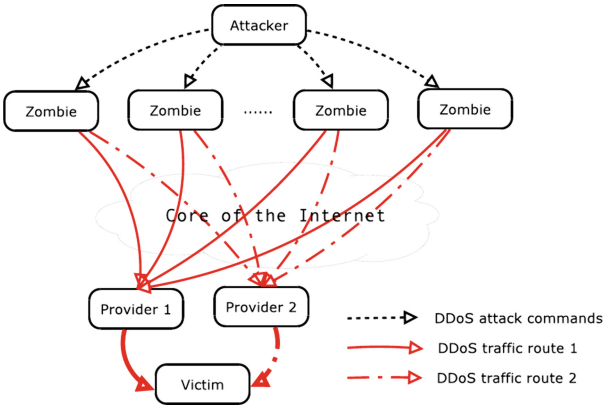

Fig. 1. An example of DDoS attack with different routes to reach the victim.

and also to which neighbors it should further propagate the reachability of the prefix. This process results in a fully-connected Internet.

DDoS Attack. Figure 1 shows a simple example of a DDoS attack where the attack traffic can have different routes to reach the victim, depending on the provider AS that the victim chooses. If the victim chooses Provider 1 as its provider AS, it will announce its prefixes via Provider 1, and as a result, the rest of the Internet (including the DDoS attackers) can reach the victim via Provider 1.

We model the DDoS attacks as follows. We assume that DDoS attacks can originate from any AS $v \in V$ on the Internet and can happen at any time. We define the set of the attack source ASes as $SRC = \{v_1^s, v_2^s, .., v_n^s\}$, where v_i^s is the ith attack source AS, and $SRC \subset V$. A DDoS attacker controls large botnets to launch the attacks, thus the size of the set SRC is usually very large. Also, we assume that any AS on the Internet can be a victim of DDoS attacks.

DDoS Defense. There are multiple ways to defend against DDoS attacks. The defense can happen at the victim side or at any third-party AS that defend for the victim, i.e., edge defense; it can also be carried out by multiple ASes on the paths of the attack traffic, i.e., in-network defense. In this paper, we allow each AS on the Internet to freely choose whether or not to participate in DDoS defense efforts, assuming that the means of communicating defense details is available to all ASes.

3.2 A Game Theoretical Model of Provider Selection

As the Internet infrastructure grows, it is common for a customer AS to have multiple provider ASes to select from. Since a customer AS can select which

providers to use independently, provider ASes compete with each other to attract more customers and thus more profit. While a key selection criterion is a potential provider's traffic forwarding service, the effectiveness of a potential provider's protection against DDoS attacks is also an important factor. A common practice is for a provider AS to charge a customer AS for 95% of the peak traffic volume (i.e., 95th percentile bandwidth metering). If a provider AS does not provide DDoS defense to filter the DDoS traffic toward the customer AS, because DDoS attacks are frequent and can bring a large amount of DDoS traffic to a customer AS, the customer AS will have to carry not only the normal traffic but also a large amount of DDoS traffic at a significantly higher cost. A rational customer AS would choose a provider AS that provides services at a lower cost, taking into account the cost of forwarding or filtering DDoS traffic by the provider AS.

Provider ASes are driven by the profit they can make while competing with other potential provider ASes for the same customer. When a provider AS invests in DDoS defense, it affects not only its own profit, but also the choice of provider AS by its current and potential customers, as well as the profit of its competitors. A provider AS that offers DDoS defense is more likely to win a customer than a provider AS that does not, and thus carrying that customer's traffic. Similarly, a provider AS that charges less for DDoS defense is more likely to attract a customer AS than a provider AS that charges more. The losing AS would then miss out on the potential profit that could have been made from the customer AS.

Below, we mathematically describe the probability that an AS will win customers from a game theory perspective. Note that a provider AS does not necessarily filter DDoS traffic for free for its customers. In this paper, we assume that each provider charges its customers for both forwarding traffic and filtering unwanted traffic.

Without losing generality, consider an example where two provider ASes v_1 and v_2 compete for a customer AS v_c. When under DDoS attack, the customer AS v_c receives a total of $T_{c,ddos}$ DDoS traffic and $T_{c,normal}$ normal traffic. Providers v_1 and v_2 must each make a decision about whether to provide DDoS filtering service and how much to charge for filtering DDoS traffic. A provider charges its customer for forwarding traffic at the rate of $r_{forward}$ and for filtering DDoS traffic—if it decides to participate in defense—at the rate of r_{filter}. If v_1 decides to participate in defense, it would charge the customer

$$r_{1,c} = r_{1,forward} * T_{c,normal} + r_{1,filter} * T_{c,ddos}.$$

In this case, it charges separately for the forwarding of normal traffic and the filtering of DDoS traffic at different rates. On the other hand, if provider v_2 decides not to participate in DDoS defense, it would charge the customer

$$r_{2,c} = r_{2,forward} * (T_{c,normal} + T_{c,ddos}).$$

Here, v_2 charges for forwarding both normal and DDoS traffic at the same forwarding rate.

The customer AS v_c chooses v_1 or v_2 as his provider based on their charges (i.e., $r_{1,c}$ vs. $r_{2,c}$). Obviously, the customer would choose the provider with the lower charge. We assume that the charging rate for filtering DDoS traffic is generally lower than that for forwarding DDoS traffic, so a provider with DDoS defense service is generally preferred over a provider without. However, the customer is likely to make errors in evaluating $r_{1,c}$ and $r_{2,c}$ due to imperfect estimation of its normal and DDoS traffic. This is called the *bounded rationality* assumption, which is widely used in the recent economics literature to capture the empirical fact that decision makers are not necessarily perfectly rational (see, for example, McKelvey and Palfrey [24] and Goeree et al. [14]). Given the errors and misinformation, the customer makes a probabilistic rather than a deterministic choice between the two providers, with a higher probability of choosing the provider with a lower expected charge.

Suppose that the customer's learned charge from provider v_1 is given by $r_{1,c} + \epsilon_1$, and that from provider v_2 is given by $r_{2,c} + \epsilon_2$, where ϵ_1 and ϵ_2 are two independent noise terms. Also define x as $r_{2,c} - r_{1,c}$. The probability that the customer will choose v_1 over v_2 as its provider, denoted as $Prob_1(c)$, is then

$$Prob_1(c) = Prob(r_{1,c} + \epsilon_1 \leq r_{2,c} + \epsilon_2) \tag{1}$$

i.e.,

$$Prob_1(c) = Prob(\epsilon_1 - \epsilon_2 \leq r_{2,c} - r_{1,c}) = Prob(\epsilon_1 - \epsilon_2 \leq x) \tag{2}$$

A common assumption is that these noise terms are extreme value distributed (usually double exponential) (see Brock and Durlauf [8,9], Durlauf and Ioannides [11], Blume et al. [5], among many others). This assumption implies that the difference of the two noise terms is logistically distributed (see McFadden [23], Anderson et al. [2], Blume [4], Brock [7] for discussions of the importance of logistic models in economics). So we have

$$Prob_1(c) = Prob(\epsilon_1 - \epsilon_2 \leq x) = \frac{1}{1 + e^{-\lambda x}}, \tag{3}$$

where λ is a parameter that measures the "noisiness" of the two noise terms. As λ increases, the customer's learned values are more accurate. Also, from Eq. (3), we can see that with a larger value of x, i.e., more savings if v_1 is chosen instead of v_2, v_1 has a higher chance of being chosen. (When x is 0, $r_{1,c}$ and $r_{2,c}$ are equal and $Prob_1(c)$ will be simply 0.5).

3.3 Profit Calculation

We assume that each provider AS tries to maximize its profit by taking into account the competition described above. As an example, consider two ASes: v_p and v_c. v_c will make a decision whether to use v_p as its provider based on the estimated charge $r_{p,c}$. The total amount of normal traffic that v_c needs a provider to carry is given by $T_{c,normal}$, and the total amount of DDoS traffic that

v_c faces is given by $T_{c,ddos}$. If v_p is chosen as the provider (wins the competition), the expected profit of v_p made from v_c is given by

$$Profit_{p,c} = Prob_p(c) \times r_{p,c}, \tag{4}$$

which is equal to the probability that v_p will be chosen by v_c as the provider times the charge of provider v_p. Note that increasing the charge $r_{p,c}$ has two opposite effects: it increases v_p's profit from handling traffic for v_c if it wins and is selected as a provider, but it also decreases v_p's probability of winning the competition due to a higher charge.

The total expected profit of an AS v_p is the sum of all the profits it can earn from its customers,

$$Profit_p = \sum_{c \in C_p} Profit_{p,c}, \tag{5}$$

where C_p is the set of all potential customer ASes of v_p.

3.4 Cost of Defense

Providing DDoS defense also incurs costs for a provider AS. We define the total defense cost function for an AS v_i as

$$Cost_i = Cost_{equip}(T_{i,ddos}) + Cost_{labor}(|C_i|),$$

where $Cost_{equip}(T_{i,ddos})$ is the equipment cost to filter the DDoS traffic toward all of its customers for an amount $T_{i,ddos}$, and $Cost_{labor}(|C_i|)$ is the labor cost to maintain all of its customers. For simplicity, we assume that the equipment and labor cost functions are the same for all ASes. We assume that $Cost_{equip}$ increases as the total traffic to be filtered increases, and C_{labor} increases as a provider AS has more customer ASes.

4 Simulation Design

In the previous section, we described our modeling of the incentives of ASes to deploy DDoS defense. To study the real-world implications of the model, we design a simulation system that allows us to simulate the DDoS defense decisions of ASes on the Internet and explore the outcomes. Designing and implementing a simulation system with more than 60,000 interconnected entities is not a trivial task. In this section, we describe our design of the simulation system and explain how the system can help us explore the outcomes of defense decisions.

4.1 Simulation Setup

Customer-Provider Pairs. The simulation relies on the full Internet topology to study all possible interactions among the provider ASes and their potential customers on the Internet. The topology includes not only the current links

(relationships) between ASes, but also all possible/potential relationships that could be established. Ideally, we would like to know which neighboring (directly connected) ASes each AS has and could have, including how it relates to them. However, relationship information is considered private and is often concealed by ASes, even for currently established ones. To best estimate the current and potential relationships between ASes, we use CAIDA AS relationship data [10, 21] from the past 10 years to compile a relatively comprehensive set of interconnection information for the entire Internet (Table 1).

Table 1. Simulation parameters and their value ranges.

parameter	meaning	value range
`ddos_ratio`	Ratio of DDoS to normal traffic volume	0.0–5.0
`do_defense`	Whether an AS participates in DDoS defense	`true`/`false`
`filter_charge`	If `do_defense` is true, how much it charges for filtering traffic	0.0–1.0

Traffic Estimation. We also use full routing tables of all collectors from RIPE RIS [27] and RouteViews [33] to construct a best-effort Internet topology. The connected peer routers are considered traffic originators in the simulation for both DDoS traffic and normal traffic. From this, we estimate the relative amount of traffic for each AS and use this as the base unit for the profit calculation (see Sect. 3).

Provider AS's Action Options. Each provider AS has several options to optimize its profit. It can choose not to participate in DDoS defense and thus charge the same rate for forwarding both normal and DDoS traffic. It can also choose to participate in DDoS defense and charge forwarding normal traffic at the regular rate, and charge filtering of DDoS traffic at a fraction of the regular rate, ranging from 0.0 to 1.0. If the filtering charge rate is 0.0, it indicates that the AS provide DDoS defense service for free; if the filtering charge rate at 1.0, the provider charges as much for filtering traffic as for forwarding it; any filtering charge above 1.0 would make filtering more expensive than forwarding for the customer, and we do not consider these cases in this paper.

DDoS Traffic Ratio. Another important factor that could affect providers' defense decisions is the severity of the DDoS attacks their customers are experiencing. We define the term *ddos_ratio* as the ratio of DDoS attack traffic to normal legitimate traffic to/from a customer AS. In this study, we range the DDoS ratio from 0.5, where the volume of DDoS traffic is half that of normal traffic, to 5.0, where the DDoS traffic is five times that of normal traffic.

Algorithm 1: Static simulation algorithm.

```
V = {v_i | v_i is an AS} ⟵ set of all ASes;
A = {a_i | a_i is a defense option} ⟵ set of defense options;
D = θ ⟵ set of defense decisions;
# first-one-to-deploy scenario
for v_i ∈ V do
    # calculate options
    for a_k ∈ A do
        for v_j is a customer of v_i do
            set competitors to no defense
            calculate probability Prob_i(j, a_k)
        end
        calculate profit ∑ Profit_{i,j}
    end
end
# last-one-to-deploy scenario
for v_i ∈ V do
    # calculate options
    for a_k ∈ A do
        for v_j is a customer of v_i do
            set competitors to defend at fixed charge
            calculate probability Prob_i(j, a_k)
        end
        calculate profit ∑ Profit_{i,j}
    end
end
```

4.2 Static Simulation

In a static simulation, each provider AS makes a defense decision based on its competitors' initial states. Once the decision is made, no further adjustment is considered. Specifically, we study the profit changes for each provider AS in the Internet when it switches from not defending to defending. We consider two scenarios in this study: the first AS to participate in defense among competitors (or *first-one-to-deploy*); the last AS to participate in defense among competitors (or *last-one-to-deploy*). The first scenario shows the likelihood of early adoption of defense solutions, while the second scenario shows how defense decisions can influence competitors' defense decisions. For each provider AS, we exhaust the defense options from no defense to defense at different charge rates and calculate the expected profit for each defense option under both scenarios. In each scenario, we also examine the total number of providers who can make more profit at a certain charge rate than by not participating in the defense.

Algorithm 2: Dynamic simulation algorithm.

$V = \{v_i | v_i \ is \ an \ AS\}$ ⟵ set of all ASes;
$A = \{a_i | a_i \ is \ a \ defense \ option\}$ ⟵ set of defense options;
$D = \{a_i | a_i \in A \ and \ v_i \in V\}$ ⟵ set of defense decisions;
$D' = \theta$ ⟵ set of previous decisions;
$isConverged$ =`false` ⟵ convergence indicator;
while *!isConverged* **do**
  `# update defense decisions`
  for $v_i \in V$ **do**
    `# calculate options`
    for v_j *is a customer of* v_i **do**
      for $a_k \in A$ **do**
        `calculate probability` $Prob_i(j, a_k)$
      end
    end
    `# find best action`
    $D_i = a_k \ where \ \max_{a_k \in A} Profit_i$
  end
  `# test convergence`
  if $D == D'$ **then**
    $isConverged =$ `true`
  else
    $D' \leftarrow D$
    $isConverged =$ `false`
  end
end

4.3 Dynamic Simulation

The simulation runs several rounds of decision making for all ASes. In each round, each AS will choose a DDoS defense effort that can optimize its profit based on the knowledge of its competitors' defense efforts from the previous round. We call such a decision rule used by the ASes the *myopic best response rule*, because each AS chooses the optimal effort level without considering the strategic changes of its competitors.

Given a configuration, we want to use the simulation to find the Nash equilibrium of the game. The equilibrium (or convergence) state represents the state where no individual player wants to unilaterally change his effort level. The detailed steps are as follows.

We initiate all ASes by setting their defense preference to no defense. As described in Sect. 3, each provider AS calculates its probability of being selected by its potential customers for each action choice it has. Using the customer-winning probabilities, the simulation system recalculates the expected profit of a provider AS, given that the AS would choose the option that can maximize the winning probability. After all ASes have updated their decisions, the simulation determines if it has converged, i.e., if any AS has made different decisions compared to the previous round. If the simulation has not converged, it continues

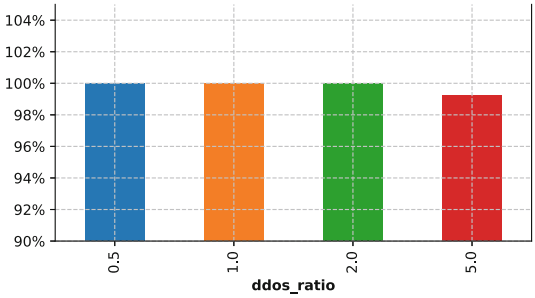

(a) Percentage of providers who can make a profit.

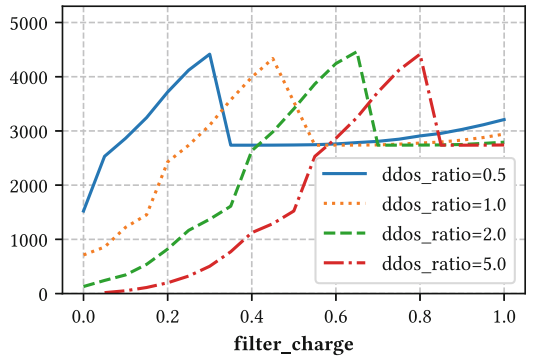

(b) Number of profitable providers at different charge rates.

Fig. 2. Profitable *first-one-to-deploy* providers.

with the previous procedure and updates the defense efforts for all ASes. If the simulation has converged, it then produces the report of the final states for all ASes. See Algorithm 2 for more details.

5 Simulation Results

In this section, we discuss the simulation results. As discussed in Sect. 4, we divide the simulation into two types:

- *Static simulation*, where only the AS under study can change its defense configuration, while other ASes, especially competitors, remain static;
- *dynamic simulation*, where each AS can change its configuration to seek higher profit.

The static simulation shows a snapshot of ASes' responses under a fixed configuration, while the dynamic simulation shows how ASes can update their defense decisions and whether the simulation achieves equilibrium. We also study these

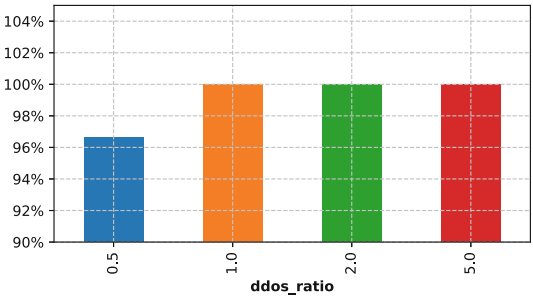

(a) Percentage of providers who can make a profit.

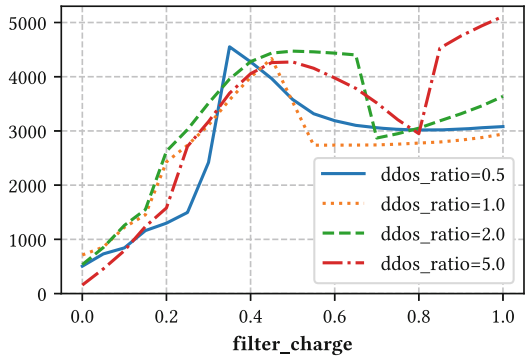

(b) Number of profitable providers at different charge rates.

Fig. 3. Profitable *last-one-to-deploy* providers.

scenarios with different DDoS-to-normal traffic ratios (*ddos_ratio*), ranging from 0.5 (i.e., the volume of DDoS traffic's is half that of normal traffic) to 5.0 (i.e., the volume of DDoS traffic is five times that of normal traffic). Each provider AS that chooses to defend can also select a different charge rate for processing and filtering DDoS traffic ($filter\_charge$).

5.1 Static Simulation

We first examine the incentives of provider ASes participating in in-network DDoS defense. Specifically, we study the profit changes for each provider AS on the Internet when it switches from not defending to defending. We consider two scenarios in this study: the first AS among competitors to participate in defense (or *first-one-to-deploy*); the last AS among competitors to participate in defense (or *last-one-to-deploy*). The first scenario reveals the likelihood of early adoption of defense solutions, while the second scenario reveals how defense decisions can influence competitors' defense decisions.

Figure 2 shows the results for the *first-one-to-deploy* scenario. We first examine how many providers can make a positive profit by switching from no defense to defense. Figure 2a shows the percentage of all provider ASes that can gain profit by switching to defense at a certain filter charge. It indicates that almost all provider ASes can make extra profit at some filter charge. Figure 2b further reveals at what charge the majority of provider ASes can make a profit by switching. It clearly shows that the number of profitable providers peaks at different filter charges; and as the DDoS becomes more severe, the peak of the filter charge increases.

Figure 3 shows the results for the *last-one-to-deploy* scenario. Figure 3a shows that almost all provider ASes can make profit at some filter charge, but there are less profitable providers when the DDoS volume is low (i.e., $ddos\_ratio = 0.5$). We further investigate at what charge most of the provider ASes can gain profit by switching when all competitors defend with 0.5 filter charge. Figure 3b shows that most of the providers can make a profit by charging slightly less than their competitors (i.e., 0.4 as opposed to 0.5). Unlike the previous scenario, it also shows that the severity of the DDoS attacks (i.e., $ddos\_ratio$) does not significantly affect the profitability of an AS when its competitors all defend.

From both results, we can see that when given the freedom to compensate DDoS defense costs by charging for filtering efforts, the majority of the providers on the Internet can find some charge rate that allows them to make profits by providing filtering services. If a provider is the first to provide services, the amount of DDoS traffic ratio affects how much it should charge to maximize profits; on the other hand, if a provider is the last to join the defense, no matter how severe the DDoS attacks are, the profitability is significantly decided by its competitor's choice. In most cases, charging similar or slightly less than the competitors would result in the best profits for the majority of the providers. In other words, charging too much would risk losing the customers altogether, while charging too little would cause the provider to miss out on a large portion of the profit it could make.

5.2 Individual Provider Profit Patterns

The previous study focused on the overall statistics of the providers that can gain profit by switching to participate in defense. We now study how each individual AS's profit can change when different $filter\_charge$ is selected. Figure 4 shows four different types of profit patterns:

- Bell-shaped profit curve with gain;
- Bell-shaped profit curve without gain;
- Increasing profit curve;
- Decreasing profit curve.

We will discuss these types of profit patterns and their indications in this section.

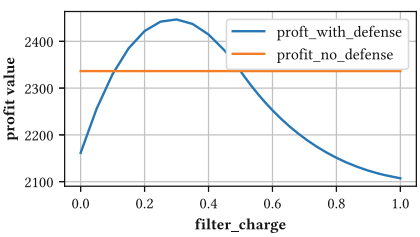

(a) Bell-shape profit curve (AS37468, 1.0 DDoS ratio).

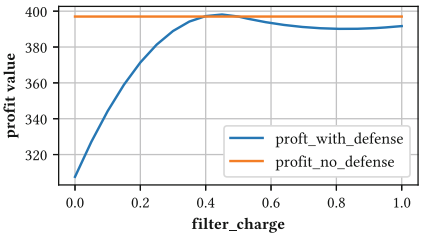

(b) Bell-shape profit curve but unprofitable (AS25227, 1.0 DDoS ratio).

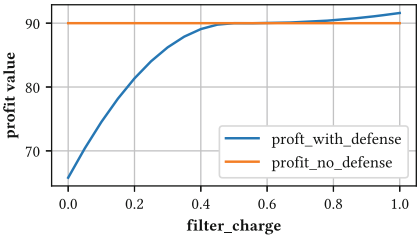

(c) Increasing profit curve (AS46455, 1.0 DDoS ratio).

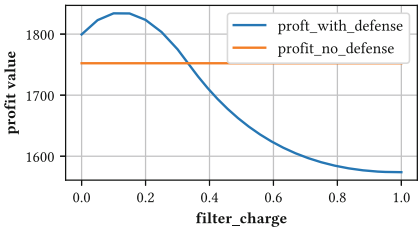

(d) Decreasing profit curve (AS37468, 0.5 DDoS ratio).

Fig. 4. Number of providers that gain by switching from not defending to defending.

Bell-Shaped Profit Curve with Gain. Figure 4a shows AS37468's profit value as the *filter_charge* changes when it participates in defense. Compared to the profit baseline when it does not participate in defense, the bell-shaped profit value curve exceeds the baseline between 0.1 and 0.5 and peaks at 0.3. The AS makes more profit by increasing the filtering charge when the charge is low (<0.3), and less profit when the charge is high (>0.3). This is a clear example of diminishing returns [30], and indicates that the AS has an incentive to participate in the defense when the charge is properly set.

Bell-Shaped Profit Curve Without Gain. Figure 4b shows a similar profit pattern with diminishing returns, but the peak of the profit when participate in defense is lower than the baseline. This figure indicates that the AS in question cannot make enough profit to justify switching to DDoS defense, regardless of the *filter_charge* choices.

Increasing Profit Curve. The increasing profit curve (Fig. 4c) indicates that the ASes have not yet reached their peak profit, even when charging at the 1.0 rate. This pattern shows that these ASes face less competitors that compete by charging prices (such as provider ASes that have customers that have no other potential providers).

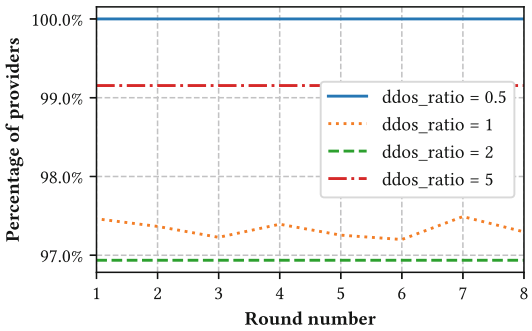

(a) Percentage of provider ASes that participate defense.

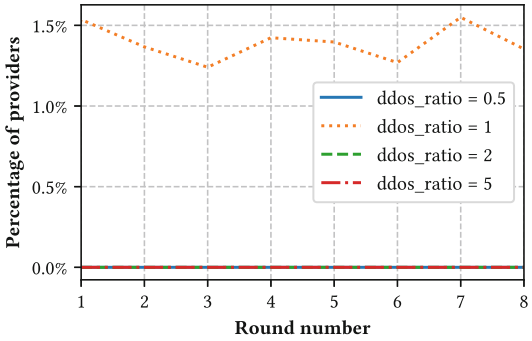

(b) Percentage of provider ASes update charges.

Fig. 5. Profitable providers when *no competitors* participate in defense.

Decreasing Profit Curve. The decreasing profit curve (Fig. 4d, after 0.1), on the other hand, indicates that the ASes have passed their peak profit at low or almost no charge for filtering DDoS traffic. Such ASes tend to be highly competitive, where charging high prices for defense would significantly decrease the chances of being selected by their potential customers.

5.3 Dynamic Simulation

We further study the decisions of provider ASes regarding DDoS defense in a more dynamic environment. In this section, we examine the results for dynamic simulation, where each AS makes a decision dynamically based on the decisions of its competitors, and the process is repeated until the decisions converge.

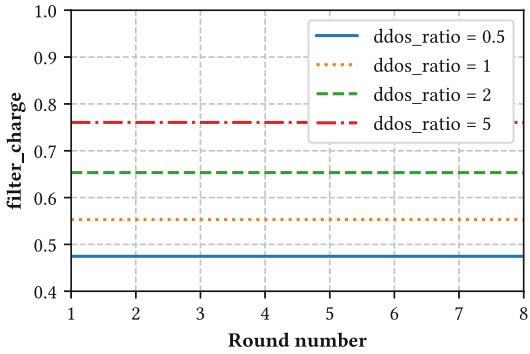

Fig. 6. Average charges.

Provider AS's Choices. In each round, a provider AS can in general decide whether or not to participate in DDoS defense based on its overall profit calculation. When deciding its options, it calculates the profit *as if*

- it does not participate in the defense;
- or it participates in defense and charges $filter\_charge$ for processing and filtering DDoS traffic.

The rate ranges from 0.0 to 1.0 with 0.1 as the increment. The AS will then choose the best option that maximizes its profit based on the calculation introduced in Sect. 3.3. Note that the current configuration of its competitors (i.e., the results of their previous decision) is taken into account during the profit calculation, and its decision will then also affect its competitors' future decisions.

Percentage of Providers Defending. We first examine the number of provider ASes that decide to participate in defense, given that each provider AS tries to maximize its profit in each round. Figure 5 shows the summary results for the dynamic simulation in terms of provider participation and configuration changes. Figure 5a shows the percentage of provider ASes that decided to participate in defense in each round of the simulation. It is clear that a very high number of provider ASes decided to participate in the defense at the very first round, and the numbers for different DDoS attack scenarios remain high and stable. This indicates that in terms of defense participation, the simulation converges quickly and result in a high level of participation for all provider ASes.

Since almost all ASes participate in defense, do they alternate their defense charges? Fig. 5b shows that there is a very number of ASes that update their $filter\_charge$ configuration and only appear in one of the simulations where $ddos\_ratio = 1$. Combining this result with the results of Fig. 5a, we can conclude that if given the opportunity to freely change and optimize their defense decisions, **the majority of the provider ASes would choose to defend and settle on their** $filter\_charge$ **rates**.

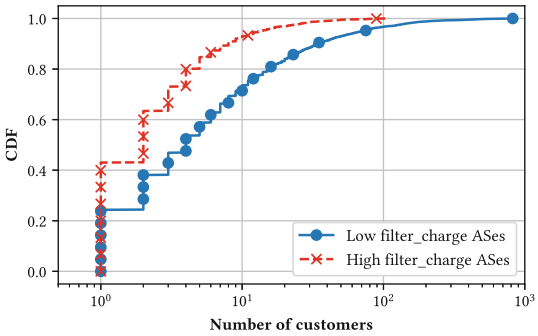

(a) CDF of number of customers.

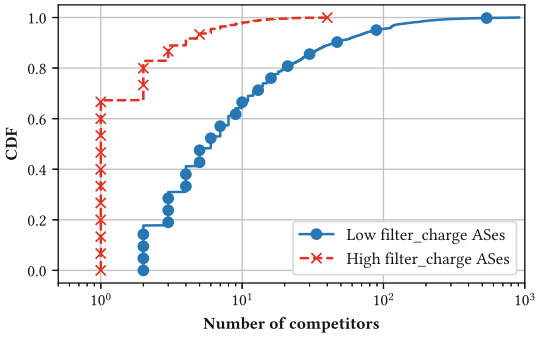

(b) CDF of number of competitors.

Fig. 7. CDF of number of customers and competitors for provider ASes.

Filter Charges. Since the majority of the provider ASes would choose to defend, the next question is how much would they charge their customers to maximize their profit? To answer this question, we dig deeper into the simulation results and examine the optimal charge of each individual AS and the overall distribution of charge values.

Figure 6 shows the average *filter_charge* for all provider ASes that decide to participate in defense under different DDoS traffic ratios. When the DDoS attack traffic is relatively small (i.e., <0.5), the average charge for filtering traffic is around 0.5, meaning that a provider AS charges its customer about half the price for filtering DDoS traffic than forwarding normal traffic. As the volume of DDoS traffic increases, the charge also increases, up to about 0.8 when DDoS attack traffic is five times the normal traffic.

We also look at the distribution of *filter_charge* for all provider ASes under different severities of DDoS attacks. Figure 8 shows the histograms of the number of provider ASes with different *filter_charge* after the simulation converges. As the DDoS ratio increases from 0.5 (Fig. 8a), 1.0 (Fig. 8b), 2.0 (Fig. 8c), to 5.0 (Fig. 8d), there is 1) a group of ASes with lower charges that move to higher

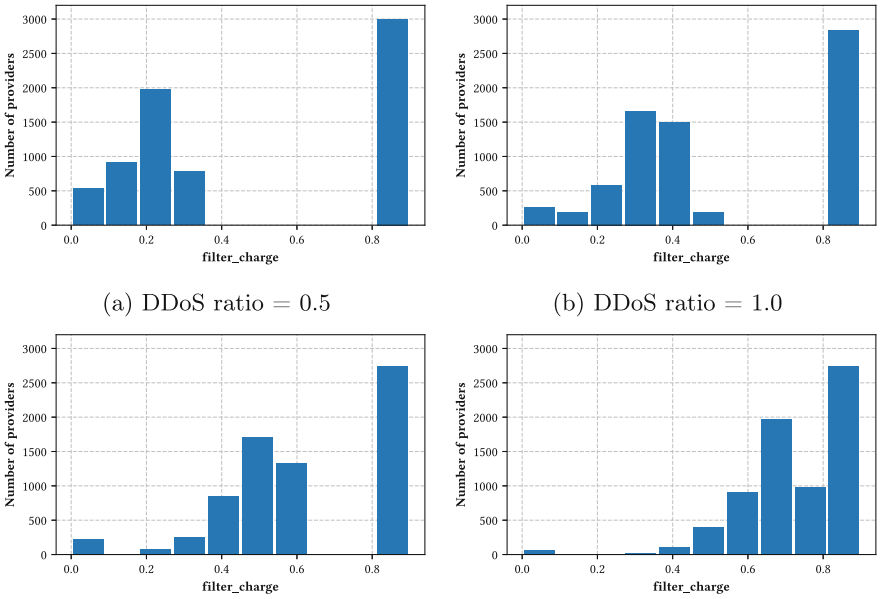

(a) DDoS ratio = 0.5

(b) DDoS ratio = 1.0

(c) DDoS ratio = 2.0

(d) DDoS ratio = 5.0

Fig. 8. Provider ASes charge distribution.

charges (which also form a bell-shape in the figures), and 2) a constant number of ASes (around 2,500 to 3,000) that always charge the highest possible rate. To understand why there are two groups of ASes with different charge patterns, we examine each AS in each group to reveal their internal correlations. Specifically, we take Fig. 8a as an example and divide all provider ASes into two groups by their $filter\_charge$: low $filter\_charge$ group where $filter\_charge < 0.5$ and high $filter\_charge$ group where $filter\_charge \geq 0.5$. Figure 7 shows the CDF plots of the number of customers and competitors for each group. It is clear that ASes with higher $filter\_charge$ have fewer customers (Fig. 7a) and fewer competitors (Fig. 7b).

5.4 Summary

In this section, we investigate the incentive of provider ASes on the Internet to participate in DDoS defense by examining the profit each provider can make under different environments, while customer ASes are free to choose the provider they want to use.

From the static simulation, we learned that most provider ASes can make profit by offering DDoS defense service to potential customers; providers reach their maximum expected profit with different charges for filtering traffic, which is affected by the severity of DDoS attacks as well as their competitors' defense decisions.

Further dynamic simulation showed that for most provider ASes, if they choose to charge for DDoS filtering that maximizes their profit (assuming no competitors offer similar services), they can all reach their stable peak profit and achieve a global stable state.

We also found that the number of competitors/customers has a strong impact on how much they should charge for DDoS filtering to achieve maximum profit. An AS can make more profit by charging more for DDoS filtering if it has weak competitions, while an AS with strong competitions needs to charge less for DDoS filtering in order to attract customers and make more profit.

6 Conclusion

In this study, we proposed a game-theoretic model that examines the incentives of ASes to invest in DDoS defense. Based on the model, we built a large-scale simulation system to examine the effects of a provider AS's topological location, level of competition, and the amount of DDoS traffic on its DDoS defense decision. From the simulation results, we observe the following patterns. The majority of provider ASes on the Internet can benefit from providing DDoS defense services to their customers if they can recover the cost of defense by charging for filtering DDoS traffic. The severity of DDoS attacks affects the rate that a provider can charge its potential customers; if a provider sees a higher volume of DDoS traffic coming toward its potential customers, it would charge more to achieve its peak profit. The level of competition also affects the rate: a provider with little competition can charge a high rate and still be profitable; a provider facing strong competition must charge less to attract customers and make a profit. These observations provide confidence that if collaborative in-network defense mechanisms mature enough, provider ASes on the Internet would have an incentive to participate in DDoS defense. We believe that such observations can further help researchers develop better strategies for designing and deploying DDoS defense solutions.

Acknowledgment. This project is in part the result of funding provided by the Science and Technology Directorate of the United States Department of Homeland Security under contract number D15PC00204. The views and conclusions contained herein are those of the authors and should not be interpreted necessarily representing the official policies or endorsements, either expressed or implied, of the Department of Homeland Security or the US Government.

References

1. Cisco annual internet report (2018–2023) white paper (2020). https://www.cisco.com/c/en/us/solutions/executive-perspectives/annual-internet-report
2. Anderson, S.P., De Palma, A., Thisse, J.F.: Discrete Choice Theory of Product Differentiation. MIT Press, Cambridge (1992)

3. Bedi, H.S., Roy, S., Shiva, S.: Game theory-based defense mechanisms against DDoS attacks on TCP/TCP-friendly flows. In: 2011 IEEE Symposium on Computational Intelligence in Cyber Security (CICS), pp. 129–136 (2011)
4. Blume, L.E.: The statistical mechanics of strategic interaction. Games Econom. Behav. **5**(3), 387–424 (1993)
5. Blume, L.E., Brock, W.A., Durlauf, S.N., Jayaraman, R.: Linear social interactions models. J. Polit. Econ. **123**(2), 444–496 (2015)
6. Bohawek, S., Hespanha, J.P., Lee, J., Lim, C., Obraczka, K.: Game theoretic stochastic routing for fault tolerance and security in computer networks. IEEE Trans. Parallel Distrib. Syst. **18**(9), 1227–1240 (2007)
7. Brock, W.A.: Pathways to randomness in the economy: emergent nonlinearity and chaos in economics and finance. Estud. Econ. 3–55 (1993)
8. Brock, W.A., Durlauf, S.N.: Discrete choice with social interactions. Rev. Econ. Stud. **68**(2), 235–260 (2001)
9. Brock, W.A., Durlauf, S.N.: A multinomial-choice model of neighborhood effects. Am. Econ. Rev. **92**(2), 298–303 (2002)
10. CAIDA: As relationships dataset (2019). http://www.caida.org/data/as-relationships/
11. Durlauf, S.N., Ioannides, Y.M.: Social interactions. Annu. Rev. Econ. **2**(1), 451–478 (2010)
12. Gao, L.: On inferring autonomous system relationships in the Internet. IEEE/ACM Trans. Netw. **9**(6), 733–745 (2001)
13. Gill, P., Schapira, M., Goldberg, S.: Let the market drive deployment: a strategy for transitioning to BGP security. In: Proceedings of the ACM SIGCOMM 2011 Conference on SIGCOMM, pp. 14–25 (2011)
14. Goeree, J.K., Holt, C.A., Palfrey, T.R.: Regular quantal response equilibrium. Exp. Econ. **8**(4), 347–367 (2005)
15. Grossklags, J., Christin, N., Chuang, J.: Security and insurance management in networks with heterogeneous agents. In: Proceedings of the 9th ACM Conference on Electronic Commerce, pp. 160–169 (2008)
16. Grossklags, J., Christin, N., Chuang, J.: Secure or insure?: a game-theoretic analysis of information security games. In: Proceeding of the 17th International Conference on World Wide Web, vol. 7, no. 1, pp. 209–218 (2008)
17. Huang, Y., Geng, X., Whinston, A.B.: Defeating DDoS attacks by fixing the incentive chain. ACM Trans. Internet Technol. **7**(1), 5–es (2007)
18. Kang, M.S., Lee, S.B., Gligor, V.D.: The crossfire attack. In: 2013 IEEE Symposium on Security and Privacy (SP), pp. 127–141. IEEE (2013)
19. Kottler, S.: February 28th DDoS incident report. GitHub Engineering (2018). https://githubengineering.com/ddos-incident-report/
20. Laszka, A., Felegyhazi, M., Buttyan, L.: A survey of interdependent information security games. ACM Comput. Surv. **47**(2), 1–38 (2014)
21. Luckie, M., Huffaker, B., Dhamdhere, A., Giotsas, V., et al.: As relationships, customer cones, and validation. In: Proceedings of the 2013 Conference on Internet Measurement Conference, pp. 243–256. ACM (2013)
22. Manshaei, M.H., Zhu, Q., Alpcan, T., Bacşar, T., Hubaux, J.P.: Game theory meets network security and privacy, vol. 45 (2013)
23. McFadden, D.: Conditional logit analysis of qualitative choice behavior. In: Zarembka (ed.) Frontiers in Econometrics, pp. 105–142 (1973)
24. McKelvey, R.D., Palfrey, T.R.: Quantal response equilibria for normal form games. Games Econom. Behav. **10**(1), 6–38 (1995)

25. Miura-Ko, R.A., Yolken, B., Mitchell, J., Bambos, N.: Security decision-making among interdependent organizations. In: 2008 21st IEEE Computer Security Foundations Symposium, pp. 66–80 (2008)
26. Papadimitriou, C.H.: Algorithms, games, and the internet. In: Proceedings of the Thirtythird Annual ACM Symposium on Theory of Computing, pp. 749–753 (2001)
27. RIPE RIS: RIPE RIS raw data (2019). https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data
28. Roy, S., Ellis, C., Shiva, S., Dasgupta, D., Shandilya, V., Wu, Q.: A survey of game theory as applied to network security. In: 2010 43rd Hawaii International Conference on System Sciences, pp. 1–10 (2010)
29. Sami, R., Katabi, D., Faratin, P., Wroclawski, J.: Practice and theory of incentives in networked systems (PINS): workshop report (2004)
30. Samuelson, P.A., Nordhaus, W.D.: Microeconomics, ISE editions (2001)
31. Shen, Y., Yan, Z., Kantola, R.: Game theoretical analysis of the acceptance of global trust management for unwanted traffic control. In: 2013 IEEE 10th International Conference on High Performance Computing and Communications & 2013 IEEE International Conference on Embedded and Ubiquitous Computing, pp. 935–942 (2013)
32. Shiva, S., Roy, S., Dasgupta, D.: Game theory for cyber security. In: Proceedings of the Sixth Annual Workshop on Cyber Security and Information Intelligence Research, vol. 34 (2010)
33. University of Oregon: Route Views project (2019). http://www.routeviews.org
34. Wu, Q., Shiva, S., Roy, S., Ellis, C., Datla, V.: On modeling and simulation of game theory-based defense mechanisms against DoS and DDoS attacks. In: Proceedings of the 2010 Spring Simulation Multiconference on - SpringSim 2010, p. 10 (2010)