Acta Informatica Pragensia
2025, Volume 14, Issue 1, pp. 1–25
https://doi.org/10.18267/j.aip.248

PRAGUE UNIVERSITY OF ECONOMICS AND BUSINESS

Article Open Access

Cloud Survivability Scenarios Under Attacks With and Without Countermeasures

Rachid Beghdad [1,2], **Faiza Benmenzer** [1,2], **Alaa Eddine Khalfoune** [1,2]

[1] Département d'Informatique, Faculté des Sciences Exactes, Université de Bejaia, Bejaia, Algeria
[2] Laboratoire LAMIE, Faculté des Mathématiques et d'Informatique, Université de Batna 2, Batna, Algeria

Corresponding author: Rachid Beghdad (rachid.beghdad@univ-bejaia.dz)

Editorial Record

First submission received:
April 16, 2024

Revisions received:
June 28, 2024
September 5, 2024
September 19, 2024

Accepted for publication:
September 27, 2024

Academic Editor:
Kang Leng Chiew
Universiti Malaysia Sarawak, Malaysia

This article was accepted for publication by the Academic Editor upon evaluation of the reviewers' comments.

How to cite this article:
Beghdad, R., Benmenzer, F., & Khalfoune, A. E. (2025). Cloud Survivability Scenarios Under Attacks With and Without Countermeasures. *Acta Informatica Pragensia*, 14(1), 1–25. https://doi.org/10.18267/j.aip.248

Abstract

Background: Despite its increasing importance, cloud computing is vulnerable to Distributed Denial of Service (DDoS) attacks, affecting data centre availability and functionality. Unfortunately, the impact of these attacks on cloud survivability remains underexplored. Most works overlook long-term resilience and lack comprehensive metrics, in-depth simulation, large-scale experiments, and combined attack and defence scope.

Objective: This study investigates the survivability of cloud environments under DDoS attacks in extreme cases, involving intensive attacks leading to cloud failure. By simulating worst-case scenarios, including thousands of attacks on large-scale clouds with and without countermeasures, we assess cloud resilience and identify the limitations of existing defences.

Methods: We conduct extensive simulations using NetLogo, modelling a cloud environment subjected to SYN flood, smurf, UDP flood, HTTP flood and malformed packet attacks. We evaluated the impact of attacks individually and in combinations, both with and without countermeasures. Each simulation involves request exchanges between end user nodes and data centres using an appropriate algorithm. We varied parameters like the number of data centres, malicious nodes, and the types and rate of attacks.

Results: The study analyses cloud resilience in terms of message delivery, available data centres, and functional node ratios, as well as tolerance and breakage thresholds. Findings indicate that cloud systems can tolerate a certain level of DDoS attack density where data centres remain accessible even without countermeasures. However, the latter greatly enhances cloud security, although their performance may decrease dramatically under extreme conditions. This highlights the importance of optimizing countermeasures, especially to handle high-intensity attacks.

Conclusion: This study provides valuable insights for cloud managers to enhance resilience and face sophisticated DDoS attacks. While current countermeasures offer initial mitigation, they are insufficient against complex and combined threats. Thus, future research should focus on developing robust, multi-layered defence mechanisms and providing data centre duplication to ensure service availability.

Index Terms

Cloud environment; Cloud survivability; Distributed denial of service attacks; DDoS; Countermeasure; Attack tolerance; Data centres.

1 INTRODUCTION

Access to data resources, including hardware and software, became easier thanks to the offered services by cloud computing systems (Ahmad et al., 2022). However, despite its advantages, these systems might be targeted by various attacks, which makes their availability questionable (Agrawal & Tapaswi, 2019; Gupta & Badve, 2017).

The availability of resources is among the most important concerns for users, as mentioned in (Varma & Reddy, 2021). Cloud systems are vulnerable to different attacks such as malware injections, man-in-the-middle attacks, side-channel attacks and distributed denial-of-service (DDoS) attacks, the latter represents a real danger on the cloud systems availability (Balobaid et al., 2016). DDoS attacks may cause damage to the overall structure and service intrusions (Agrawal & Tapaswi, 2019; Varma & Reddy, 2021).

The ability of cloud systems to maintain functionality under attacks is called “survivability” or “attack tolerance”; hence, assessing the survivability of cloud computing is very important to determine how long these systems can keep working under different attack scenarios.

Taking into consideration the amount of research on cloud systems DDoS attacks, there are still some shortcomings. Generally, the previous research in the literature covered only performance metrics, ignoring the survivability of cloud systems. In addition, most previous studies simulated only small-scale networks, which is not always the case in cloud systems; hence, generalizing the findings on real wide-scale networks is questionable.

Furthermore, despite the huge effects of cloud systems attacks on the interconnection between users and data centres, only a few studies have covered these effects. Moreover, the impact of combined attacks and incorporating multiple defence solutions are frequently overlooked in the literature, which leaves an open gap. Moreover, existing research tends to avoid pushing simulations to the point of total cloud structure failure, which prevents a comprehensive understanding of the impacts of attacks and cloud tolerance in extreme cases.

To address these shortcomings, this study aims to investigate the following research questions: How resilient is cloud computing under various realistic DDoS attack scenarios, and what are the limitations of current countermeasures in ensuring cloud survivability? Specifically, the study focuses on worst-case attack scenarios. In worst-case attack scenarios, our research aims to study cloud tolerance by simulating dialogues occurring between an end user and a data centre under hundreds and thousands of attacks, respectively, without countermeasures, until the total destruction of data centres and the whole 1000-node and 10,000-node clouds.

Furthermore, this study highlights the limitations of some proposed solutions in extreme cases. For instance, we demonstrate that the cloud can still operate by simulating a 10,000-node cloud without countermeasures and subjecting it to hundreds of attacks; nevertheless, if existing countermeasures are compromised, this can lead to the crushing of the cloud.

The remainder of this paper is structured as follows: Section 2 reviews related work in the field. Section 3 provides an overview of the attacks addressed in our study. In Section 4, we present our simulation-based study and analysis. Finally, Section 5 concludes the paper and discusses potential directions for future research.

2 LITERATURE REVIEW

During our selection process, we concentrated on finding and reviewing scholarly papers from peer-reviewed journals and conference proceedings. The search was conducted between February 2023 and May 2024, primarily using the Scopus citation database to gather relevant literature. We used search terms like "DDoS attacks," "cloud computing," "attack tolerance in a cloud," and "cloud survivability." Our review focused on articles published from 2016 to 2024 to include recent advancements and significant studies in cloud security.

DDoS attacks on cloud infrastructure have gained significant research attention, yet gaps remain, particularly in the areas of data centre availability, large-scale experiments, diverse attack types and mitigation strategies. Balobaid et al. (2016) investigated various DoS and DDoS attacks on cloud environments, where mitigation strategies focused on software and hardware firewalls along with IDS/IPS. The research demonstrated that although the firewalls were found useful in reducing the intensity of such attacks, they could not mitigate high-volume DDoS attacks due to evolving attack tools and Trojans. However, their simulation had a very narrow scope, with no metrics for evaluating data centre availability.

Alosaimi et al. (2016) studied the impact of economic denial of sustainability on cloud services and presented an improved DDoS defence. They performed OPNET simulations to assess different scenarios. Their studies showed that the firewall successfully prevented DDoS assaults. However, the study focused solely on HTTP traffic and ignored other forms of DDoS assaults.

Zhao (2017) suggested DPDK-based DDoS detection and defence technology, using the BP algorithm to identify anomalous traffic and DPDK to forward regular packets. They demonstrated the effectiveness of the system by generating SYN flood packets with Mpktgen and directing them towards a victim's computer. However, the experiment was limited in scale, with few virtual machines, and the paper did not discuss the system's availability, indicating a need for further research.

Corrêa et al. (2021) used machine learning with cloud/fog telemetry data to detect DDoS attacks by monitoring resources such as memory, CPU, disk and bandwidth. Experiments on OpenStack showed high accuracy in detecting SYN flood and GET flood attacks. However, the study was limited to small-scale OpenStack experiments, tested only two types of DDoS attacks and did not include any investigations of mitigation mechanisms.

Similarly, Wani et al. (2019) studied DDoS attack identification in cloud environments using random forest, naïve Bayes, support vector machine and decision tree algorithms. They evaluated the accuracy, precision and recall of DDoS and normal packets. However, the study used a single attack tool and a limited IDS-generated dataset, it was also limited to specific algorithms and ignored mitigation methods.

Xu et al. (2021) conducted a study analysing DoS threats to cloud-based multi-robot systems. The authors proposed three novel attack techniques that would exploit several aspects of the system, such as the network, micro-architecture, and function parameters. They conducted evaluations and case studies to prove how far attacks can cause serious degradation in both performance and safety. However, the experiments were performed in a controlled environment, limiting their real-world applicability. Moreover, the study does not present defence solutions for enhancing resilience in cloud-robotic systems against attacks.

Ferretti et al. (2021) presented a new architectural paradigm for a cloud computing environment. The proposed system focuses on the concepts of distributed trust and strong identifying threat and attack mechanisms. While it illustrates great resilience and flexibility, it faces challenges such as reliance on survivable databases and balancing request confidentiality with intrusion detection.

Liu and Xing (2021) analysed the survivability and vulnerability of a cloud RAID storage system, considering disk faults and cyber-attacks. They proposed a continuous-time Markov chains-based method for disk-level analysis and combinatorial methods for system-level analysis. The practical application of these methods is shown in a case study on a cloud RAID 5 system. The study investigated the impact of various parameters on disk and system survivability and invulnerability through numerical analysis. However, the work has limitations such as simplifying assumptions, scalability concerns, and a lack of extensive validation through real-world testing or simulations.

Inspired by predator-prey dynamics in nature, Mthunzi and Benkhelifa (2017) introduced an analogy-based approach to advance the survivability of cloud environments. They emphasise the need for a standardised definition of survivability that includes resistance to attacks, rapid recovery, intrusion detection, and adaptability to evolving threats, ensuring the availability of cloud services. However, the approach is limited by reliance on theory, practical challenges, and the need for diverse defences.

Table 1 provides a synthesis and comparison of characteristics of the aforementioned works. It highlights the focus areas, targeted attacks, mitigation strategies evaluated and the scale of experiments conducted. Dialogue analysis involves the examination of interactions between an endpoint node and a data centre until the complete disappearance of data centres. Combined attacks highlights whether combined attacks were studied. Simulation depth indicates the extent to which the simulation was conducted, including whether it considered the total breakdown of the cloud structure. Estimated metrics include various metrics that highlight cloud survivability. These metrics are essential for assessing how well cloud systems can maintain functionality under attacks. This comparison helps identify the unique contributions and limitations of each study.

Table 1. *Comparison of related works.*

Approach	Focus	Single attack (with/without countermeasures)		Combined attacks	Combined defence	Experiment scale	Dialogue analysis	Simulation depth	Estimated metrics				
		Attack studied	Mitigation studied						MDR	ADR	FNR	TT/TB	ER
Balobaid et al. (2016)	Impacts and mitigation of DoS/DDoS attacks in cloud	SYN flood UDP flood	Software-based firewalls hardware-based firewalls IDS/IPS capabilities	x	x	Single node (2VM)	x	x	x	x	x	x	x
Alosaimi et al. (2016)	Enhanced DDoS mitigation	HTTP flood	Firewall	x	x	26 nodes	x	x	✓	x	x	x	x
Zhao (2017)	DDoS prevention strategy using DPDK	SYN, HTTP flood	DPDK, BP algorithm	x	x	NS	x	x	✓	x	x	x	x
Corrêa et al. (2021)	DDoS detection using cloud telemetry	SYN, GET flood	ML algorithms: kNN, random forest	x	x	11 nodes	x	x	x	x	x	x	x
Wani et al. (2019)	Machine learning for DDoS detection	NS	ML algorithms: support vector machine, naïve Bayes, random forest	x	x	7 nodes	x	x	x	x	x	x	x
Xu et al. (2021)	DoS threats in cloud-based multi-robot systems	Network flooding, micro-architecture, contention attacks	Not addressed	x	x	3 nodes	x	x	✓	x	x	x	x
Ferretti et al. (2021)	Survivable zero trust for cloud computing	Attacks target system components for unauthorized cloud resource access	Zero trust with survivable components	✓	x	NS	x	x	x	x	x	x	x
Liu & Xing (2021)	Cloud RAID system survivability/vulnerability	Disk fault attacks	Markov chain method	x	x	NE	x	x	x	x	x	x	x
Mthunzi & Benkhelifa (2017)	Survivability in cloud computing environments	General threats and attacks	Predator-prey inspired analogies	x	x	NE	x	x	x	x	x	x	x

Notes: ADR = available data centre rate, MDR = message delivery ratio, FNR = functional node rate, TT = tolerance threshold, BT = breakage threshold, ER = error rate, NS = not specified, NE = no experiment.

Despite the increasing interest in investigating DDoS attacks in the cloud and numerous studies in the literature, significant weaknesses persist, including:

- **Lack of comprehensive metrics:** Previous studies have focused on metrics such as response time, throughput, and detection time, which, while important, do not provide a comprehensive understanding of how cloud systems maintain functionality under sustained attacks. Essential metrics for evaluating cloud

survivability, like long-term availability and resilience, have often been overlooked. These particular metrics fail to fully capture the ability of cloud systems to survive. They don't take into account the system's long-term availability or its resilience when encountering extreme conditions, factors that are vital for assessing how well a cloud system can withstand attacks and remain functional over time. This gap in the existing metrics leaves large blind spots in understanding the true resiliency of cloud systems. Our study bridges this gap by integrating long-term impact metrics that give full insight into the functionality of cloud systems against sustained attacks, including message delivery ratio, available data centre rate, functional node rate, tolerance threshold, and breakage threshold.

- **Gap in simulation depth and dialogue analysis:** Previous studies have not extensively simulated interactions between endpoint nodes and data centres until the complete disappearance of data centres, nor have they pushed simulations to the point of total cloud structure failure. This gap in research leaves a significant lack of understanding of the full impact of DDoS attacks on data centres and the cloud ability to withstand such high attack densities. Our study aims to fill this gap by conducting simulations that assess how long data centres can withstand attacks before total collapse and by evaluating the depth of simulation to understand cloud resilience under extreme conditions.
- **Limited simple and combined attacks:** While several studies have simulated simple attacks, only a few have investigated combined attacks. Among the cited works, only Ferretti et al. (2021) delved into the study of combined attacks focusing solely on three types. However, while Ferretti et al. (2021) explored combined attacks, their focus was narrow and the combinations studied were limited to specific scenarios. In contrast, other works have concentrated on a limited number of single attacks, highlighting a gap in understanding the broader impact of attacks on cloud survivability. Our study addresses this gap by simulating a vast number of attacks, including both simple and combined attacks, until the total destruction of the cloud.
- **Lack of combined defence mechanisms:** Most studies overlook the necessity of combined defence mechanisms to address multiple simultaneous attack vectors, thereby limiting the effectiveness of proposed solutions in real-world scenarios. To fill this gap, we use an integrated combination of protection mechanisms in a strong defence strategy against complex attack scenarios.
- **Limited experiment scale:** Most works perform small-scale experiments involving a single server or few machines, which lack the complexity and scale of the real cloud environment, resulting in incomplete evaluations of the cloud's resiliency and survivability. In this regard, our work focuses on large-scale experiments that accurately evaluate the resiliency and survivability of the cloud under extreme scenarios.

Our study aims to address these limitations by conducting extensive experiments and analysing a wider range of DDoS attacks. To the best of our knowledge, no studies have examined the impact of combining five attack types. Specifically, we assess the cloud resistance to realistic worst-case data centre attack scenarios, including multiple simultaneous attacks and identify weaknesses in current literature solutions under extreme conditions.

3 STUDIED ATTACKS AND COUNTERMEASURES

In this section, we analyse various types of DDoS attacks and their corresponding countermeasures, addressing critical gaps identified in the literature. Most previous studies focused only on limited attack scenarios, whereas large-scale and attack combinations scenarios were neglected; both are essential factors in assessing the tolerance of cloud systems. In addition, combining multiple defence mechanisms has not received deserved attention despite its importance in encountering multiple simultaneous attack vectors. By addressing these gaps, we aim to enhance the robustness of our research methodology and findings.

3.1 Studied attacks

Cloud computing enables resource sharing among multiple users, introducing security vulnerabilities, particularly to DoS and DDoS attacks. DoS attacks may decrease packet delivery rate and cause buffer overflow. According to (Mishra et al., 2021), these attacks are classified into three classes depending on their goal, destroying nodes physically, disrupting limited resources, and altering configurations. DDoS attacks are mainly designed to target service availability (Velliangiri et al., 2021). They can be application-based, targeting data centres to disrupt legitimate services, or protocol-based, seeking to overload network bandwidth (Yan et al., 2016).

The existing research in cloud security have not highlighted emerging threats and mitigation strategies properly, Attacks like malformed packets, smurf and IP spoofing remain unexplored, moreover, multi-vector attack scenarios have not been sufficiently studied. Our study addresses these shortcomings by considering a wider range, including the following attacks:

- **SYN flood attack:** According to Shah et al. (2022), these attacks comprise 84.6% of all DDoS attacks, overwhelming data centres with SYN requests and making them inaccessible to legitimate users.
- **Smurf attack:** This attack floods the victim's resources using ICMP echo-reply messages, these messages are generated from broadcast data centres (Bouyeddou et al., 2021).
- **UDP flood:** In this attack, random ports on a host are flooded, blocking legitimate user requests (Devi & Subbulakshmi, 2021; Potluri et al., 2020).
- **HTTP attack:** This attack overwhelms data centres with HTTP requests, causing a denial of service for users (Khandare et al., 2023; Varma & Reddy, 2021).
- **Malformed packet:** These attacks involve sending packets with identical source and destination IP addresses, causing system crashes (Mishra et al., 2021).

Despite rapid progress in recent years, reaching an efficient scheme to detect cloud system attacks is still a key objective, as mentioned in (Ogwara et al., 2022). Combining attacks has become more prevalent in recent times; attacks such as HTTP flood, smurf attack, SYN flood, malformed packets, and IP spoofing are being combined to exploit the cloud vulnerabilities in order to exhaust its services (Corrêa et al., 2021; Osanaiye et al., 2016). Our study examines combinations of attacks to find out how they interact to cause more disruption. This reveals the limitations of current detection systems and underscores the need for robust and multi-faceted detection mechanisms. Understanding combined attacks helps develop adaptive defences and create more resilient cloud environments.

3.2 Studied countermeasures

The DDoS detection methods in cloud environments have been studied in numerous research studies, which focused on signature-based, anomaly-based, and hybrid techniques. However, these studies are based on single-layer defences without taking into consideration layered approaches for real-time threat mitigation; according to Osanaiye et al. (2016), anomaly-based and hybrid methods are the most popular approaches for DDoS detection in cloud environments.

Our study analyses recent hybrid anomaly-based defences designed for cloud data centres; we picked up approaches that include comprehensive quantitative metrics that are assessable under extreme conditions, incorporating filters, third-party auditors and hybrid systems. We chose the countermeasures based on their ability to respond to the aforementioned threats; these countermeasures have proved their effectiveness in previous studies, such as the system developed by Devi et al. (2021), who incorporated a variety of approaches to mitigate the SYN flood and UDP flood attacks effectively. The chosen countermeasures ensure a comprehensive defence strategy, thereby enhancing overall robustness. Moreover, some of the selected countermeasures use new methods, such as clustering for anomaly detection (Raja Sree & Mary Saira Bhanu, 2020) and genetic algorithms for intrusion detection (Nsabimana et al., 2020), contributing to the advancement of cloud security research.

In contrast to previous works, we implemented and evaluated the defence schemes in wide-scale cloud environments to imitate real cloud systems, which significantly contributes to this field. In the following, the studied defence mechanisms are presented:

- **Defence 1:** Devi et al. (2021) designed a system that filters malicious traffic to mitigate SYN flood and UDP flood attacks using security enforcement policies. The system captures the spatio-temporal behaviour of security events using complex event processing and stops attack sources with a threat management system.
- **Defence 2:** Raja Sree and Mary Saira Bhanu (2020) proposed a third-party auditor to identify the origin of HTTP flood attacks. The auditor traces packets, analyses traffic patterns and generates alerts. It preprocesses and normalizes features from VM network logs and cloud data centre access logs, then uses fuzzy bat clustering to detect anomalies.
- **Defence 3:** Nsabimana et al. (2020) presented a hybrid intrusion detection and prevention system using a signature-based method and a genetic algorithm to detect and prevent smurf attacks. The cloud is divided

into clusters, each monitored by a hybrid NIDS model. Alerts are sent to a logging system and forwarded to Splunk for visualization.

- **Defence 4:** Sultana et al. (2019) proposed a defence against cloud DDoS attacks caused by IP spoofing and malformed packets. Their method filters out invalid packets using a modified hop count filtering (HCF) technique. The technique examines the SYN flag, TTL, source port and source IP of each packet.

Our study uniquely implements all these defence mechanisms across small, medium and large-scale cloud environments, significantly contributing to the field. In the following section, we assess the cloud ability to withstand various simultaneous attacks aiming to compromise data centre availability.

4 ATTACK TOLERANCE IN CLOUD ENVIRONMENT

Cloud data centres are networks of virtualized data centres that offer on-demand services such as computing power, storage and network resources. These services cater to global users who rely on the cloud for various applications (Kanniga Devi et al., 2020). Despite their advantages, cloud data centres are vulnerable to a multitude of attacks, which may lead to failures and consequently impaired service availability. Thus, attack tolerance is defined as the system resilience in the face of attacks, which is crucial for cloud data centres. Our study investigates the resilience of cloud environments under worst-case scenarios: a variety of simultaneous attacks, attacks with and without countermeasures; after that, we highlight the limitations of existing defence strategies.

To model the cloud environment, we employ a graph $G = (Ni, Ci)$, where Ni represents the nodes (users and data centres) and Ci signifies the communication links within the cloud. Figure 1 provides a visual representation of this model where the nodes with the highest degree correspond to data centres.

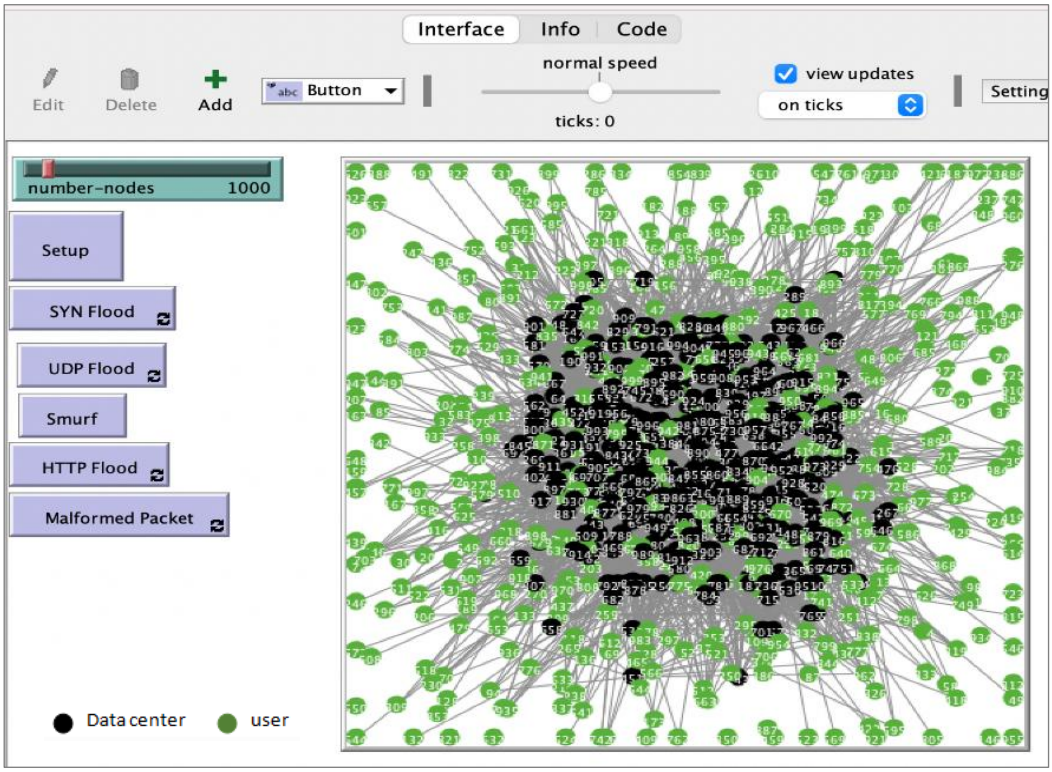

Figure 1. *Cloud environment representation, see also Appendix A.*

4.1 Simulation description

In this study, we employed NetLogo v6.2 for simulation-based analysis, given its ease of use, graphical capabilities, strong community endorsement and ability to generate and analyse large-scale simulations. The simulations were run on an Apple M1 chip computer with 8 GB of RAM.

Each simulation iteration initiates a specific number of requests that randomly originate from source nodes and target data centre nodes. The following simplified exchange algorithm is utilized to control the interaction:

```
Source node:
Random delay /*to avoid any collision*/
Send REQUEST /*Send request message to access the stored data*/
Receiver node:
      If (destination address==data centre address) then
      Send REPLY
      else
      Forward REQUEST
```

To assess attack tolerance in diverse conditions, we ran exhaustive simulations across several scenarios, including:

- a cloud with one or more attacks of the same attack type,
- a cloud subject to combined attacks,
- a cloud with attacks and countermeasures.

Metrics were averaged over multiple simulation runs to account for the inherent randomness of cloud generation. The parameters that we varied during the simulations comprised the number of data centres and other nodes, the quantity of malicious nodes, as well as the types and numbers of both attacks and countermeasures.

The NetLogo environment was initialized with a blank template, upon which we designed our cloud ecosystem. We defined the world dimensions and populated it with grid cells representing cloud nodes. Agents symbolizing key components such as data centres, users and attackers were subsequently added using NetLogo's *Breed* command. We then launched simulations to assess various DDoS attacks, incorporating an attacker agent that sent a surge of connection requests to the target data centre to render it unresponsive. The results were subsequently analysed to assess cloud resilience (see Appendix for graphical interfaces of some simulation scenarios).

In the following subsections, we offer an in-depth evaluation of the DDoS attacks and defences under study, along with their implementation specifics.

4.1.1 Attack strategies

This subsection describes how various cyber-attacks were implemented using NetLogo:

- **SYN flood attack:** An attacker agent initiates a SYN flood by inundating the data centre with SYN requests while neglecting SYN-ACK messages. We employed NetLogo to create this attacker agent, guided by pseudocode and the "*create-links-with*" command for establishing connections.
- **Smurf attack:** Using the "*nw*" extension in NetLogo, a network topology is initialized. The attacker uses the "*nw:send*" command to send ICMP echo requests with a spoofed source address.
- **UDP flood attack:** An attacker agent sends UDP packets to the target data centre using the "*nw:send*" command in NetLogo.
- **HTTP flood attack:** The attacker agent sends HTTP GET requests to the target data centre, making use of the "*nw:http-get*" command. The "*NW-HTTP extension*" is required.
- **Malformed packet attack:** The attacker uses "*nw:create-raw-packet*" and "*nw:set-packet-field*" commands in NetLogo to generate malformed packets and send them to the target data centre.

4.1.2 Defence strategies

This subsection details the implementation of various defence strategies using NetLogo. Each strategy includes specific extensions, commands, rules and procedures, described as follows:

- **Defence 1:** To implement a complex event processing system for cloud network security, we created rules to monitor traffic and detect SYN and UDP attacks using "*if*" statements in NetLogo. To prevent and mitigate these attacks, we defined a rule that checks the number of SYN packets sent to a data centre, triggering an alert or taking action to block the source IP address if it exceeds a certain threshold. Similarly, for UDP flood attacks, we set up a rule to monitor UDP traffic and detect a sudden increase in the number of packets sent to a data centre, triggering an alert or taking action to shape the traffic and block the source IP address if this

increase exceeds a certain threshold. Using NetLogo's "*ask*" and "*if*" statements, we apply these policies to different nodes and traffic flows in the cloud, creating a more secure cloud infrastructure (Figure 2).

```
to detect-syn-flood
  ask nodes [
    if syn-count > syn-threshold [
      ; Detected TCP SYN flood attack
      set alert-message (word "TCP SYN flood attack detected from IP: " ip-src)
      print alert-message
      if not member? ip-src blocked-ip-list [
        set blocked-ip-list lput ip-src blocked-ip-list
      ]
    ]
  ]
end

to detect-udp-flood-attack
  ask nodes [
    if udp-count > udp-threshold [
      ; Detected UDP flood attack
      set alert-message (word "UDP flood attack detected from IP: " ip-src)
      print alert-message
      if not member? ip-src blocked-ip-list [
        set blocked-ip-list lput ip-src blocked-ip-list
      ]
    ]
  ]
end

to apply-intrusion-prevention-policy
  ask nodes [
    if member? ip-src blocked-ip-list [
      set protocol "BLOCKED"
    ]
  ]
end
```

Figure 2. *Example of rules for defence 1 – detecting and mitigating SYN flood and UDP flood attacks.*

- **Defence 2:** We implemented fuzzy bat clustering in NetLogo by creating a list of HTTP requests for clustering. Using the distance function, we calculated similarities between requests and assigned degrees of membership with the "*fuzzify*" function. The "*defuzzify*" function assigned each request to a cluster based on membership degrees until cluster centres stabilized. To detect abnormal traffic, we set a membership threshold, identifying requests below it as outliers. These were further analysed for potential attacks. We used the "*fuzzy-and*", "*fuzzy-or*" and "*fuzzy-not*" functions in NetLogo to create rules, classifying clusters as normal or attack clusters (Figure 3).

- **Defence 3:** For smurf attack defence, we configured "*Snort-IDS*" with signature rules using the "*BehaviorSpace extension*". Deployed on each data centre node, "*Snort-IDS*" scans traffic packets using the "*foreach*" command to match known signature rules. "*Ask*" commands preprocess packets, extracting features such as packets per second and IP addresses. "*If*" statements check for smurf attack criteria. We created NetLogo procedures to analyse traffic patterns and inform nodes of potential attacks using genetic algorithms. For instance, we set up a rule that generates an alert when the number of ICMP packets from a specific source IP exceeds a threshold within a time frame. Another rule generates an alert when many ICMP packets with broadcast addresses are detected quickly. Each data centre node preprocesses incoming traffic packets, and "*Snort-IDS*" and genetic algorithms analyse the data for attack detection. Alerts are forwarded to a logging system and analysed by a rule-based expert system for severity and response (Figure 4).

```
; Implement Fuzzy Bat Clustering Algorithm
to fuzzy-bat-clustering
  ; Initialize variables
  let request-list network-traffic

  ; Calculate similarity between requests
  let similarity-matrix create-similarity-matrix request-list

  ; Fuzzify each request
  let membership-matrix create-membership-matrix similarity-matrix

  ; Initial cluster centers
  set cluster-list initialize-clusters membership-matrix

  ; Repeat clustering process until convergence
  let converged? false
  while [not converged?] [
    let new-cluster-list k-means-clustering membership-matrix cluster-list
    if cluster-list = new-cluster-list [
      set converged? true
    ]
    set cluster-list new-cluster-list
    set similarity-matrix create-similarity-matrix request-list
    set membership-matrix create-membership-matrix similarity-matrix
  ]

  ; Classify clusters
  classify-clusters cluster-list
end

; Classify Clusters
to classify-clusters [clusters]
  foreach clusters [
    c -> let req-per-sec fuzzy-and (get-membership c "requests-per-sec") (get-membership c "unique-IPs")
         let http-type fuzzy-not (get-membership c "HTTP-request-type")
         let attack-prob fuzzy-or req-per-sec http-type
         if attack-prob > 0.5 [
           set c-type "attack"
         ]
         else [
           set c-type "normal"
         ]
  ]
end
```

Figure 3. *Example of rules for defence 2 – detecting and mitigating HTTP flood attack.*

```
 to preprocess-packet [packet]
   let src-ip item 2 packet  ; Extract source IP address
   let dst-ip item 3 packet  ; Extract destination IP address
   let timestamp item 4 packet  ; Extract timestamp of packet
   let packet-size item 5 packet  ; Extract packet size
   let packets-per-second packet-size / (timestamp - last-timestamp)  ; Calculate packets per second
   set last-timestamp timestamp  ; Update last timestamp for next packet

   ; Return a list of extracted features
   report (list packets-per-second src-ip dst-ip)
 end

 to smurf-attack-detection
   let packets snort:scan-traffic
   foreach packets [
     packet ->
     let features preprocess-packet packet
     let packets-per-second first features
     let src-ip item 1 features
     let dst-ip item 2 features

     ; Rule for detecting high ICMP packet rate from a single source
     if packets-per-second > threshold [
       snort:generate-alert "High ICMP packet rate from " + src-ip
     ]

     ; Rule for detecting ICMP packets with broadcast addresses
     if dst-ip = "255.255.255.255" [
       snort:generate-alert "ICMP broadcast detected from " + src-ip
     ]
   ]

   ; Forward alerts to the logging system
   snort:forward-alerts
 end
```

Figure 4. *Example of rules for defence 3 – detecting and mitigating smurf attack.*

- **Defence 4:** To implement the modified Hop count filtering technique in NetLogo, we created a "*packet*" breed with "*hop-count*", "*source-ip*" and "*destination-ip*" attributes. We implemented filter-packet procedures in cloud nodes, which periodically remove invalid packets based on the is-valid-packet? function. The filter-packets procedure collects all packets within a radius of 1 and checks their validity using the "*nw:turtles-in-radius* " command. If a packet is found to be invalid, the procedure removes it using "*die*" primitive. The "*is-valid-packet?"* function checks if the packet hops count falls within the valid range and if the source IP address is different from the destination IP address using the "*ip extension*". If both conditions are true, the function reports true, indicating that the packet is valid. Otherwise, the function reports false, indicating that the packet is invalid (Figure 5).

```
to filter-packets
  ;; Filter out invalid packets
  let packets-to-check nw:turtles-in-radius 1
  foreach packets-to-check [
    if not is-valid-packet? self [
      ;; Remove packet if not valid
      ask self [ die ]
    ]
  ]
end

to-report is-valid-packet? [packet]
  let valid-hop-count? (packet-hop-count packet >= 1 and packet-hop-count packet <= 64)
  let valid-ip-spoof? (packet-source-ip packet != packet-destination-ip packet)

  ;; Return true if both conditions are met
  report valid-hop-count? and valid-ip-spoof?
end
```

Figure 5. *Example of rules for defence 3 – detecting malformed packet attack.*

Following the implementation of aforementioned attacks and countermeasures, we simulated various scenarios to assess cloud resilience. We recorded results, calculated averages and analysed the impact of the attacks using specified metrics.

4.2 Evaluation metrics

To evaluate the resilience and efficiency of cloud systems under various attack scenarios, we employ a range of metrics. These metrics are primarily focused on evaluating cloud availability, including the rate of data centre availability and the system tolerance to different scales of attacks. A detailed breakdown of these metrics and simulation parameters is presented in Table 2.

Table 2. *Simulation parameters.*

	Parameters	**Descriptions**
Environment	Material	Apple M1 chip with 8 GB of RAM
	Simulator	NetLogo version 6.2
	Cloud size	Varied from 50 to 10,000 nodes
	Number of data centres and users	Varied in accordance with cloud size; data centres have maximum degree in the graph
	DDoS implemented	Smurf, SYN flood, UDP flood, HTTP flood, malformed packet attacks
	Number of attackers	Varied from 0 to cloud destruction
	Defence implemented	Defences against smurf attack, SYN flood, UDP flood, HTTP flood and malformed packet attacks
Metrics	Message delivery ratio (MDR)	Percentage of successful control messages
	Available data centre rate (ADR)	Percentage of accessible data centres
	Functional node rate (FNR)	Proportion of successful user-data centre interactions

	Parameters	Descriptions
	Tolerance threshold (TT)	Number of attackers that the system can tolerate
	Breakage threshold (BT)	Number of attacks needed to compromise cloud services

4.3 Simulations and performance analysis

In this section, we assess cloud performance in large-scale settings through various simulation scenarios. The first scenario evaluates a cloud environment under different types of single attacks, comparing results without and with countermeasures. The second scenario examines the combined effect of all the studied attacks and the corresponding defences, assessing the cloud ability to withstand multifaceted threats. Performance curves for various metrics are presented to illustrate the cloud capabilities and limitations under these conditions.

4.3.1 Case with single attacks

This subsection evaluates the cloud under different single attacks, comparing results without and with countermeasures. Key performance metrics are analysed to determine the cloud tolerance and the effectiveness of defences.

Message delivery ratio

Figure 6 illustrates the impact of attacking nodes on the MDR for a 10,000-node cloud, both without and with countermeasures, across different attack types.

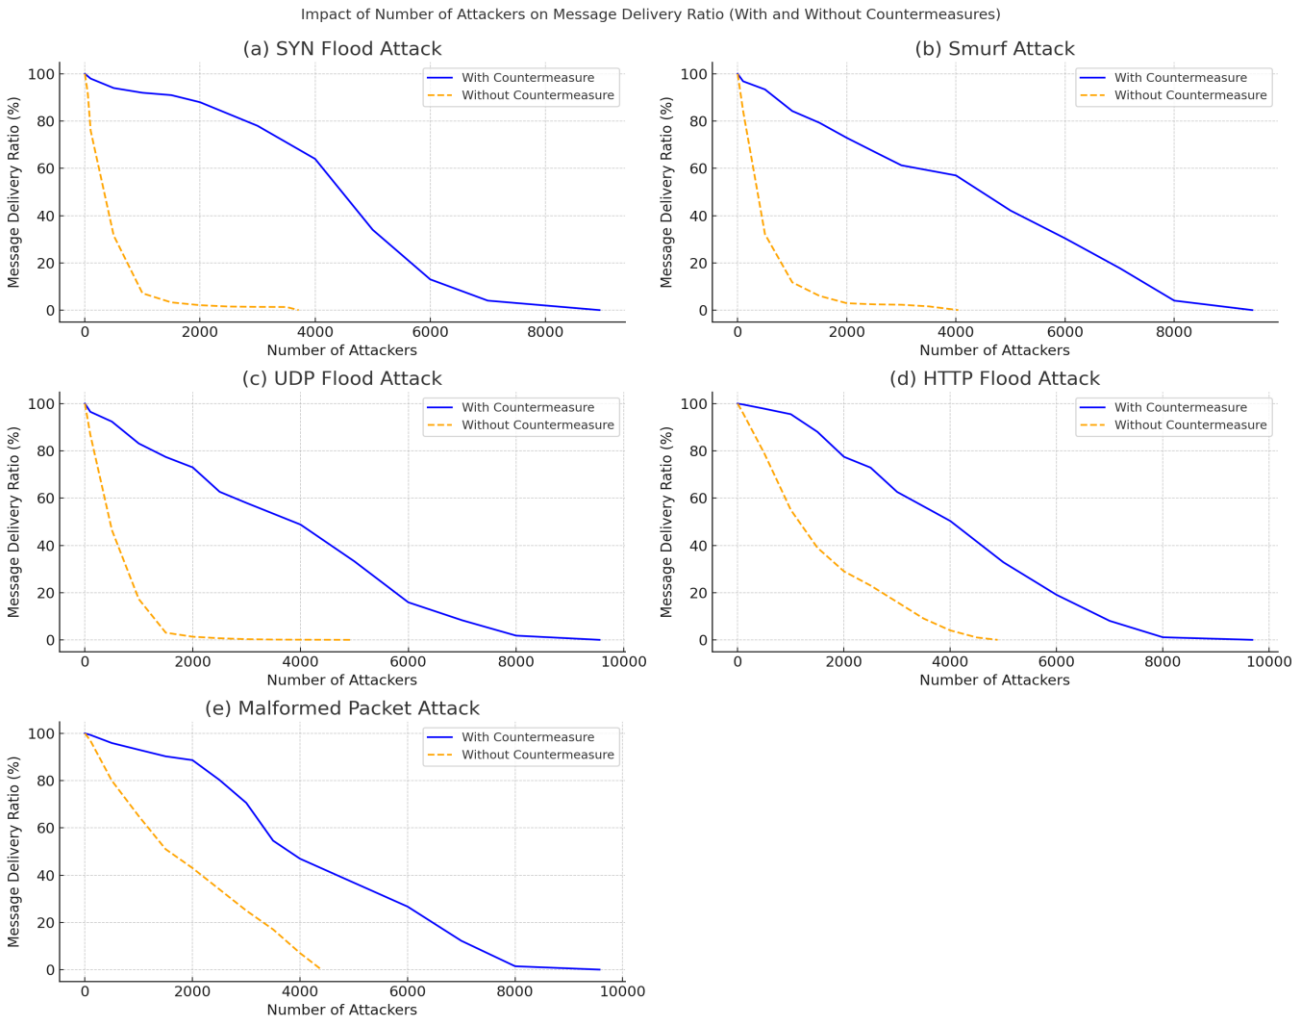

Figure 6. *Impact of number of attackers on message delivery ratio under different attacks.*

According to the findings illustrated in Figure 6, an increase in the percentage of attacking nodes causes a decrease in the MDR across all scenarios. Without countermeasures, the reduction in MDR is due to the increase in SYN, UDP, HTTP and fake echo requests, which strain the cloud resources and limit its ability to respond to legitimate user requests. Consequently, the data centre capacity to send REPLY messages is diminished, leading to a significant decrease in MDR.

Implementing countermeasures improves the MDR across all scenarios by detecting malicious nodes and blocking fake requests. The effectiveness of individual countermeasures is evident as they significantly mitigate the impact of attacks. For instance, in the SYN flood attack scenario (Figure 6a), the MDR falls dramatically as malicious nodes are presented to approach 0% with only 20% attackers without the defence mechanisms; however, introducing the countermeasure improves maintaining the MDR notably, in which it needs around 80% of attacking nodes to reach 0%.

For the smurf attack (Figure 6b) and UDP flood attack (Figure 6c), the MDR curve shows similar patterns without countermeasures: it decreases dramatically as attack density lifts to reach 0% when the number of attackers exceeds 20%, while it falls less sharply for both HTTP flood (Figure 6d) and malformed packets (Figure 6e) to reach 0% by introducing a 40% attacking density; however, when presenting the defence mechanisms, the performance improves notably and the MDR declines regularly to reach 0% as attack density touches 80%.

These MDR drops in the scenarios with countermeasures occur because the volume of malicious traffic at these thresholds overwhelms the countermeasure capacity. Each countermeasure is designed to handle a specific amount of malicious traffic. When the number of attacking nodes exceeds this capacity, the countermeasure becomes less effective and the cloud resources are consumed by processing the attack traffic, leaving insufficient resources for legitimate user requests. This also leads to many more compromised nodes.

Available data centre rate

Figure 7 shows the effect of the attacking node density on the ADR in a 10,000-node cloud, both without and with countermeasures. According to Figure 7, data centre availability in a 10,000-node cloud diminishes dramatically as the number of attackers rises, regardless of the attack type. In SYN flood, HTTP flood and UDP flood attacks, the decrease can be justified by the surge in attack requests, making many data centres unreachable. For smurf attacks, the decline is attributed to an influx of forged ICMP queries, heightening the risk of data centre crashes. Similarly, in malformed packet attacks, data centre accessibility decreases due to increased malformed IPs.

Implementing defences enhances cloud availability by identifying malicious nodes and blocking attack traffic. However, the effectiveness of countermeasures varies across attack types. In the SYN flood scenario (Figure 7a) and smurf attacks (Figure 7b), the ADR goes down gradually to reach 0% at around 90% of malicious nodes. For UDP flood (Figure 7c), HTTP flood (Figure 7d) and malformed packet attacks (Figure 7e), the countermeasures perform a little better, as the ADR needs around 95% of attack density to reach 0%.

The performance of the countermeasures decreases under high density of attacks due to their finite capacity to filter and block malicious traffic. When the number of attackers increases beyond this capacity, the data centres become overwhelmed by the volume of attack requests. This leads to a sharp decline in data centre availability, as the data centres are unable to process legitimate requests while dealing with the excessive load of attack traffic.

The existing studies have focused on individual attack types and have not considered the availability of data centres during diverse and concurrent attack scenarios. This gap emphasizes the importance of conducting more comprehensive research to fully understand the impact on data centre availability.

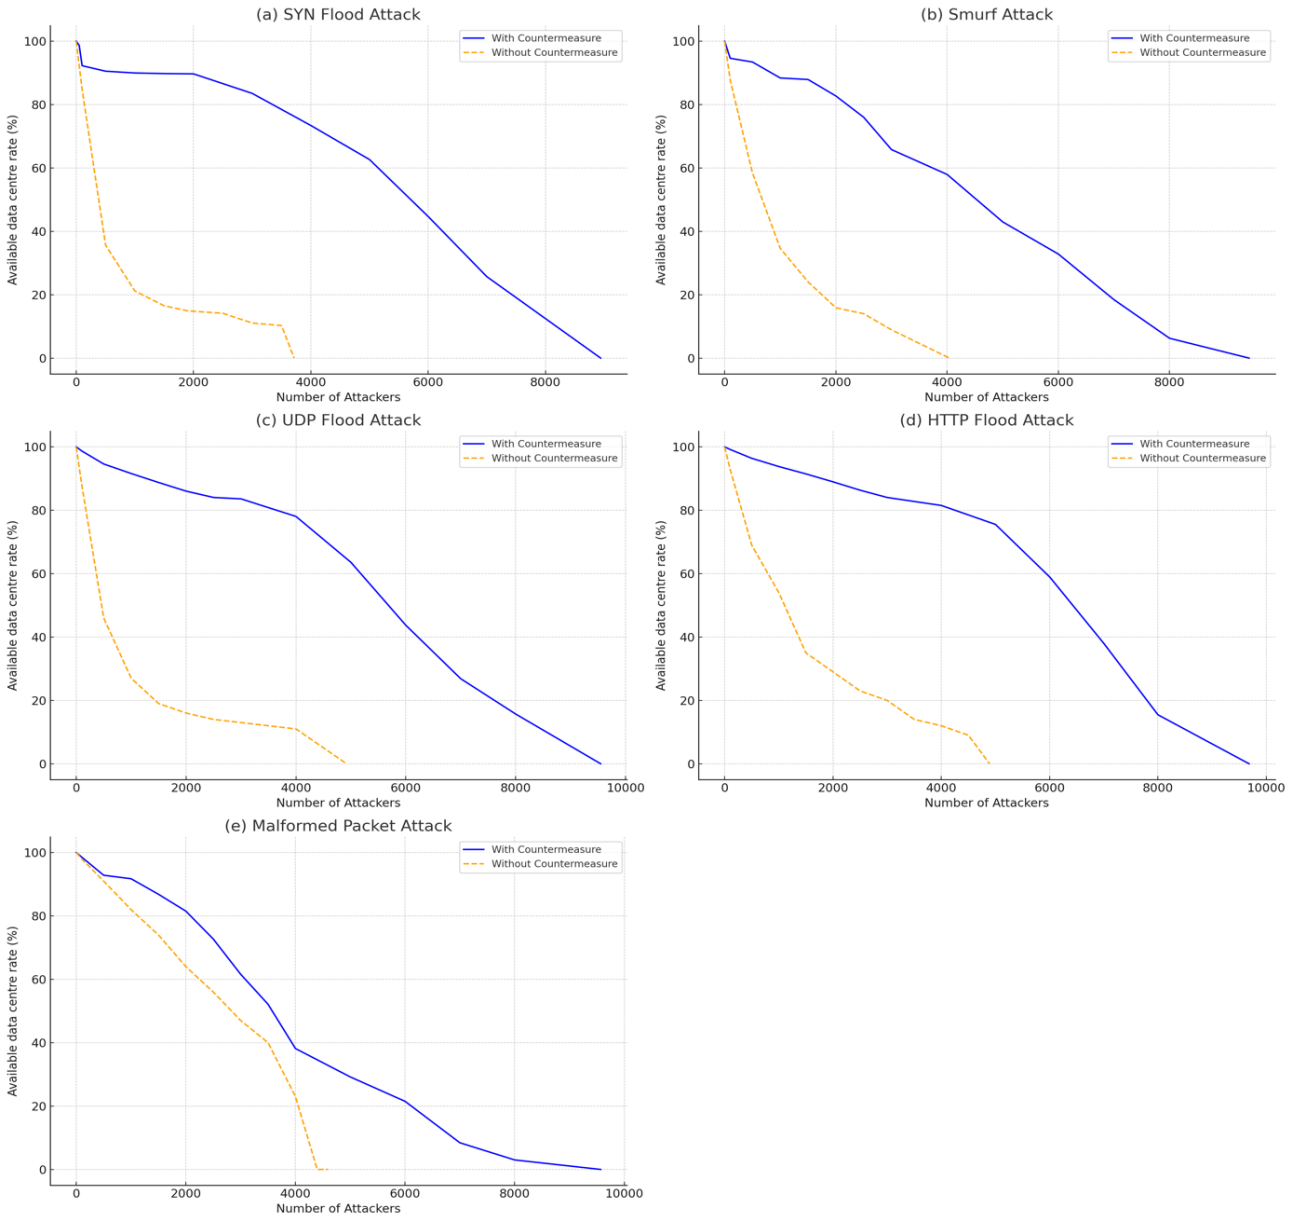

Figure 7. *Impact of number of attacking nodes on available data centre rate in 10,000-node cloud under different attacks.*

Functional node rate

Figure 8 illustrates the impact of the number of attackers on the FNR in a 10,000-node cloud, both without and with countermeasures. According to Figure 8, the rate of functional nodes falls as more attacking nodes are present in the cloud. For SYN, UDP and HTTP flood attacks, this is attributed to a surge in corresponding attack requests, disrupting user data access and decreasing the number of functional nodes. In smurf attacks, the decline occurs due to amplified fake ICMP echo requests causing denial of service and a decrease in functional nodes. In malformed packet attacks, the issue stems from a large volume of malformed IPs, preventing user access and decreasing the number of functional nodes.

Implementing countermeasures notably boosts the FNR across all scenarios. However, the FNR still declines when the number of malicious nodes grows. In the SYN flood attack scenario (Figure 8a), the FNR gradually decreases and hits 0% when the fraction of malicious nodes exceeds 88%. In the smurf scenario (Figure 8b), the FNR drops to 0% beyond 92% of malicious nodes. For UDP flood (Figure 8c), HTTP flood (Figure 8d) and malformed packet (Figure 8e) scenarios, the FNR reaches 0% when the attack density is around 95%.

The secret behind these drops is that countermeasures have a limit to the number of malicious requests they can filter out. When the number of malicious nodes exceeds this limit, the functional nodes are overwhelmed by the

attack traffic, preventing them from processing legitimate user requests and maintaining normal operations. As the attack intensity increases, the number of functional nodes declines sharply once the countermeasure capacity is surpassed.

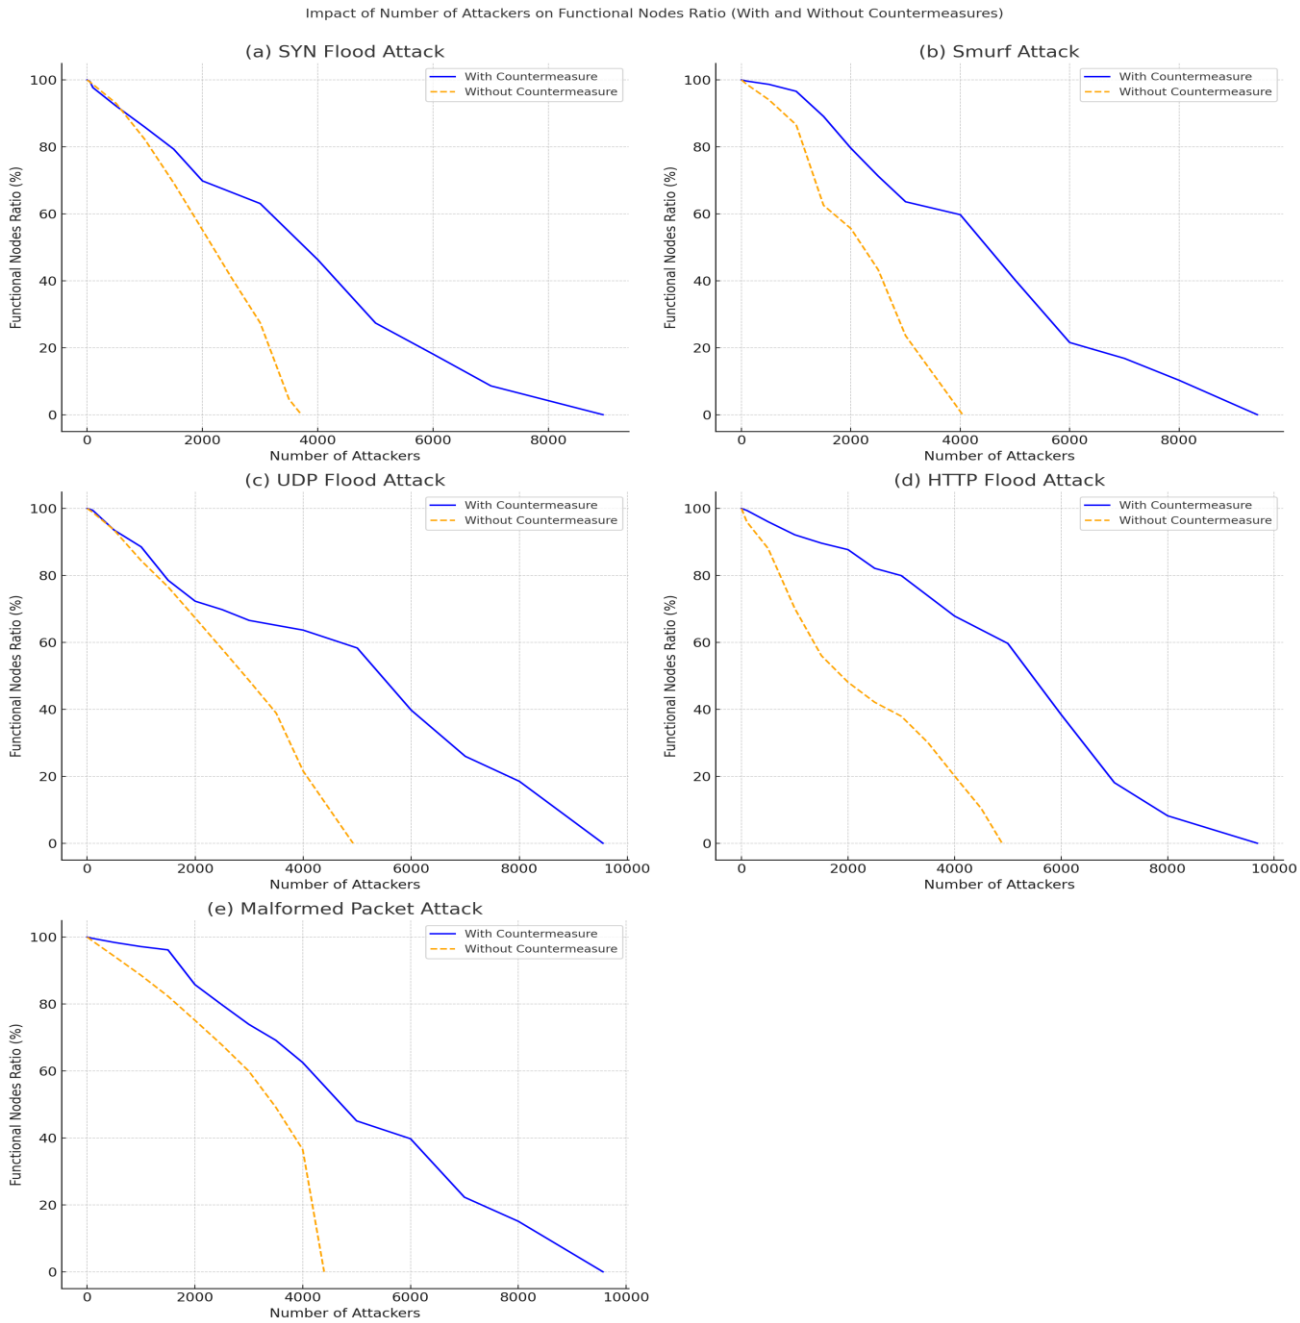

Figure 8. *Impact of number of attackers on functional node rate in 10,000-node cloud under different attacks.*

Although numerous studies have examined the impact of various attack types on cloud performance, there is a lack of comprehensive research that specifically addresses the rates of functional nodes.

Tolerance and breakage thresholds

In this subsection, we assess the cloud tolerance and breakage thresholds. Figure 9 provides a comparative analysis of the results in a 10,000-node cloud environment under various attack scenarios, both with and without countermeasures.

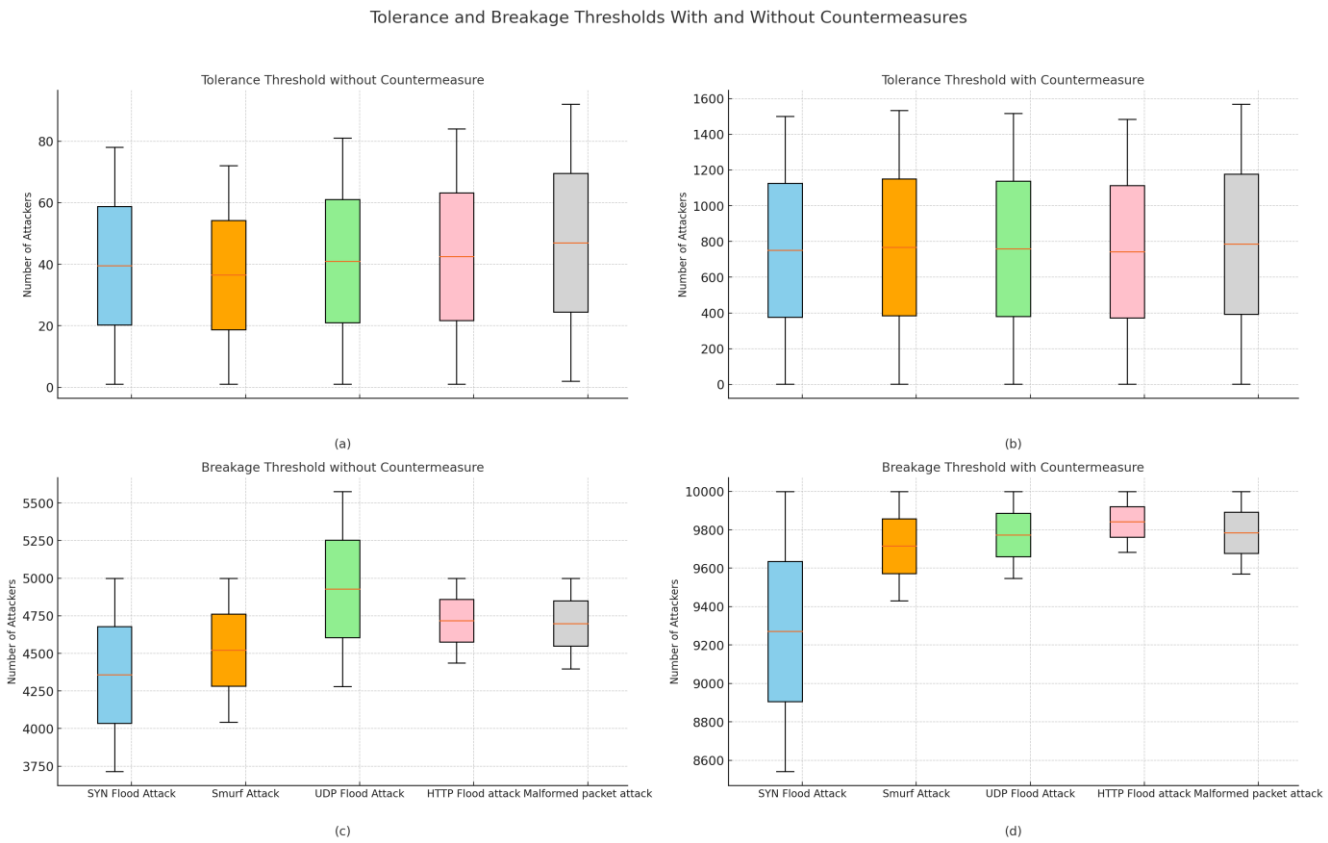

Figure 9. *Tolerance and breakage thresholds in 10,000-node cloud without countermeasures.*

According to Figure 9(a) and without countermeasures, the 10,000-node cloud can tolerate a maximum of 78, 81, 84, 92 and 72 malicious nodes for SYN flood, UDP flood, HTTP flood, malformed packet and smurf attacks, respectively. Thus, the cloud is still running even if it is attacked by 1 to 78 malicious nodes for the SYN flood, 1 to 81 malicious nodes for the UDP flood, 1 to 84 malicious nodes for the HTTP flood, 1 to 92 malicious nodes for the malformed packet attack and 1 to 72 malicious nodes for the smurf attack.

Nevertheless, the cloud is destroyed when the number of malicious nodes falls between 3715 and 5000 for the SYN flood, 4280 and 5577 for the UDP flood, 4397 and 5000 for malformed packets, 4435 and 5000 for the HTTP flood and 4043 and 5000 for smurf attacks as shown in Figure 9(c). Figure 9(b) demonstrates that countermeasures considerably enhance the 10,000-node cloud resilience against various attacks. The cloud can tolerate 1 to 1500 nodes for the SYN flood, 1 to 1516 for the UDP flood, 1 to 1483 for the HTTP flood, 1 to 1568 for malformed packets and 1 to 1534 for smurf attacks. Figure 9(d) illustrates that the number of attackers needed for a complete failure in the cloud varies from one attack to another, ranging from 8542–10,000 for the SYN flood to 9430–10,000 for the smurf attack, highlighting the effectiveness of countermeasures in mitigating risks.

Previous studies have focused on individual attack types, providing valuable insights into cloud vulnerabilities. However, these studies often overlook the complexity and interactions that arise when multiple attack types occur simultaneously. As a result, their findings may not fully capture the real-world scenarios where combined attacks can exploit cloud weaknesses in different layers and services. Additionally, no study has specifically evaluated the tolerance and breakage thresholds of clouds under various attack scenarios, highlighting a significant gap in the current literature.

4.3.2 Case with combined attacks

This subsection examines the combined effect of all the studied attacks and corresponding defences on cloud performance. It assesses the cloud ability to withstand multifaceted threats by analysing key performance metrics. Notably, there is a scarcity of studies that comprehensively evaluate data centre availability, message delivery ratios and tolerance and breakage thresholds under combined attack scenarios, which underscores the novelty and significance of this research.

The first scenario deploys a mix of the five prior attacks, each distributed over one-fifth of the cloud. The second scenario follows the same structure but incorporates countermeasures for each attack type. Key performance metrics were measured in both scenarios to assess the impact on the entire cloud. Figure 10 shows these combined attacks against data centres on five continents while processing a large number of user requests.

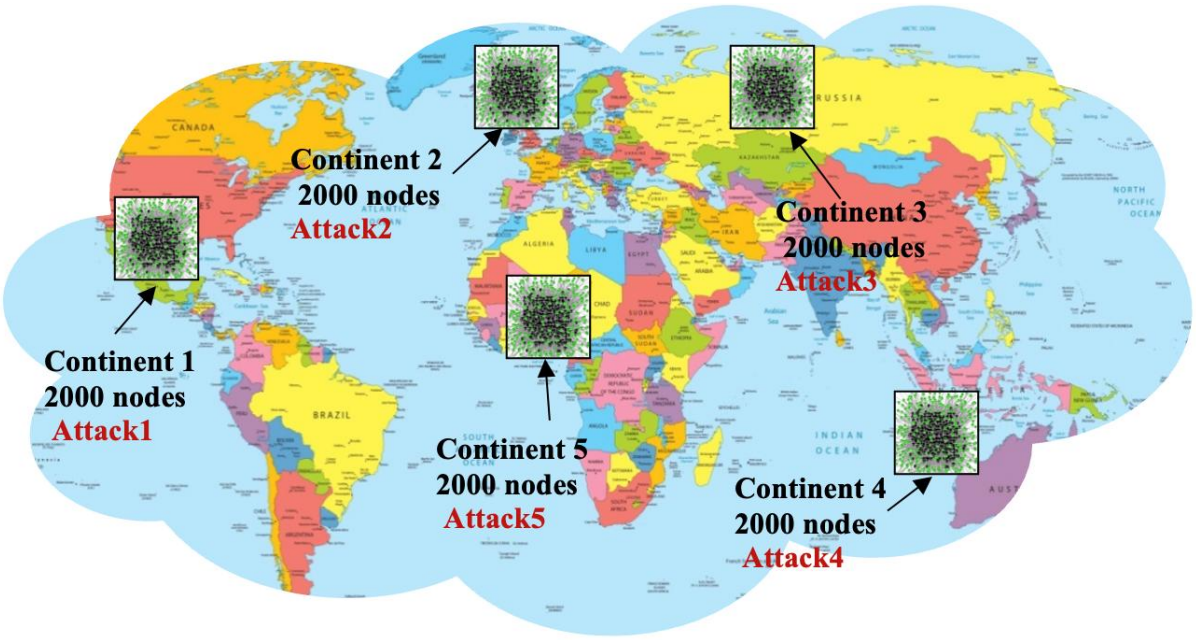

Figure 10. *Combined attack scenario in 10,000-node cloud.*

Throughout the experiment, we evaluated availability, message delivery ratios and functional node rates. Figure 11 shows how these metrics fluctuate based on the number of attackers, both with and without defences.

Message delivery ratio

According to Figure 11(a), without countermeasures, the MDR significantly drops as the number of malicious nodes increases. This decline is due to the amplified malicious traffic consuming data centre resources, which limits the cloud ability to handle legitimate user requests, leading to a higher number of compromised nodes. With countermeasures, the MDR improves substantially, although it still declines beyond certain thresholds. For instance, the MDR starts to decline sharply when malicious nodes make up over 40% of the cloud. In this case, corrupted nodes will increase, such that the MDR will reach 0 at around 90%. This shows that the defences are effective up to a point but can be overwhelmed by high-intensity attacks.

Available data centre rates

As shown in Figure 11(b), the availability of data centres diminishes sharply without countermeasures due to the surge in attack requests. Implementing defences helps maintain higher data centre availability, but similar to the MDR, the ADR begins to decline when malicious nodes exceed 50% of the cloud and reaches zero at around 90%. This indicates that while defences mitigate the impact, they can still be overwhelmed by large-scale attacks.

Functional node rate

Figure 11(c) illustrates that without defences, the rate of functional nodes decreases dramatically as attack intensity increases. The surge in attack traffic overwhelms the compromised nodes, preventing them from functioning properly. With defences in place, the FNR shows a marked improvement but also declines beyond specific thresholds. It starts to drop significantly when malicious nodes make up over 45% of the cloud and reaches zero at

around 90%. This demonstrates that the defences improve resilience but have limitations under extreme attack conditions.

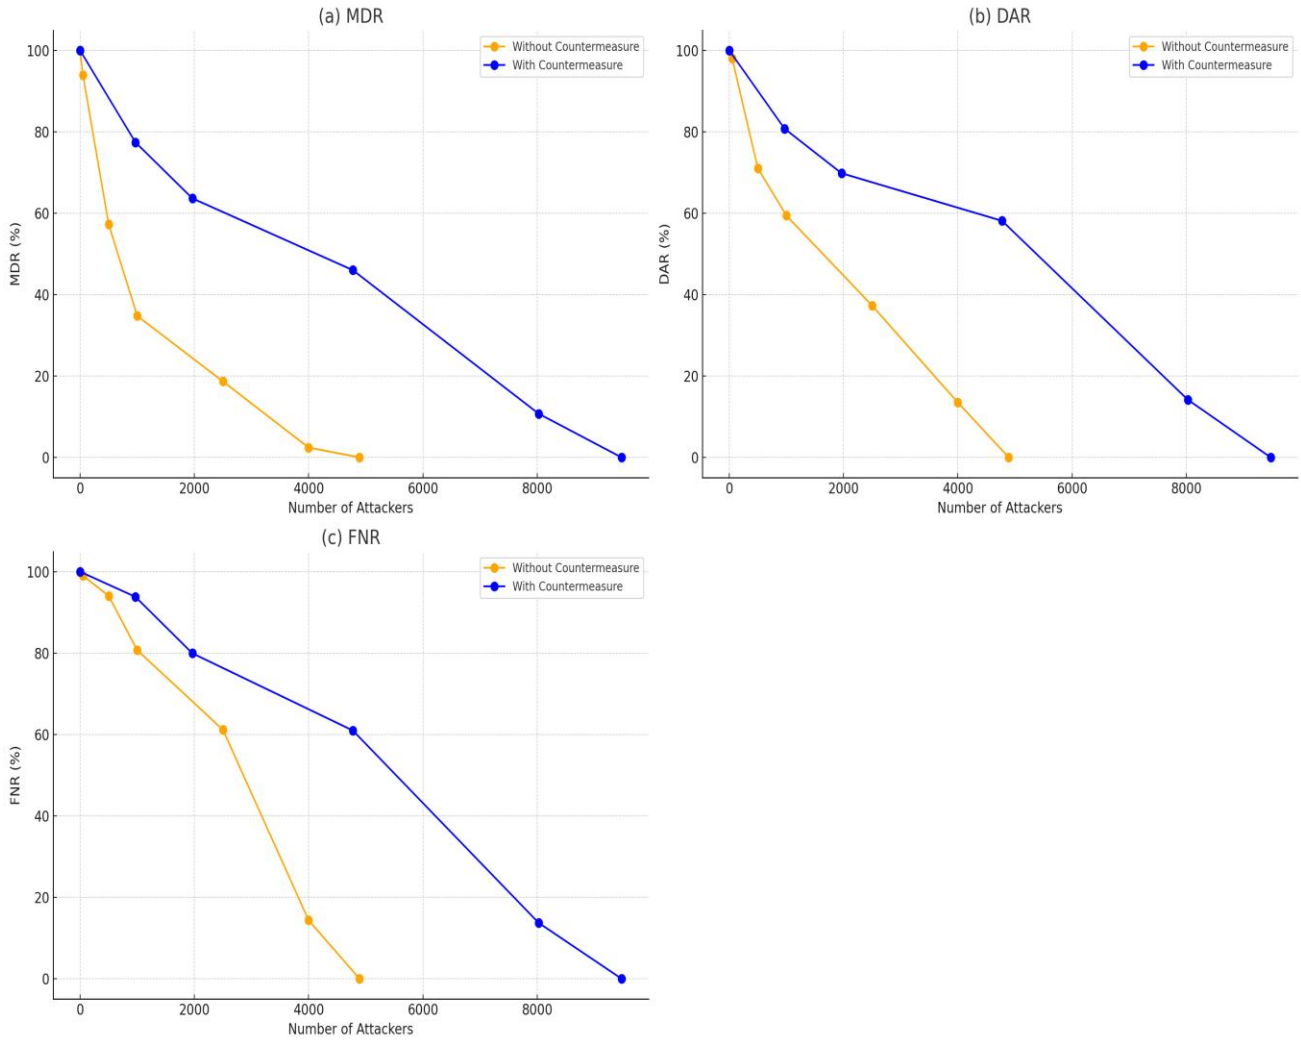

Figure 11. *Impact of combined attack density on 10,000-node cloud with and without combined countermeasures.*

Tolerance and breakage thresholds

Furthermore, we estimated the tolerance and breakage thresholds. The experiment results are illustrated in Figure 12. As shown in Figure 12, without countermeasures, the number of attackers needed to destroy a 10,000-node cloud ranges from 4890 to 5000 for mixed attacks. Conversely, the cloud can still function even if attacked by 1 to 88 malicious nodes from the combined attacks. This indicates that the 10,000-node cloud can tolerate up to 88 combined attacks distributed globally. So, the cloud is still running even if it is attacked by 1 to 88 malicious nodes combined.

With countermeasures, a 10,000-node cloud requires between 9475 and 10,000 combined attackers for complete destruction while it can remain functional under attack from 1 to 1950 malicious nodes. This indicates that with countermeasures, the 10,000-node cloud can tolerate up to 1950 combined attacks spread over different continents. This demonstrates that defences significantly extend the cloud tolerance but have a finite limit beyond which they cannot prevent failure.

Combined attacks result in higher resource depletion and greater overall impact on cloud performance compared to individual attacks. For instance, a 20% density of combined attackers decreases the ADR to around 70% as shown in Figure 11(b), which is lower than what these attacks achieve individually. The mixed attack traffic increases complexity, making it harder for defences to identify and mitigate threats effectively. Attackers often use

combinations of attacks to overwhelm cloud defences and evade detection. These combinations are common in real-world scenarios as they exploit different vulnerabilities simultaneously.

While individual countermeasures can improve key performance metrics, their effectiveness diminishes when faced with combined attacks. Countermeasures show significant improvements in metrics such as the MDR, ADR and FNR up to certain thresholds, but struggle to handle high volumes of attack traffic. This highlights the need for more robust and scalable defence mechanisms to effectively handle multifaceted threats in cloud environments.

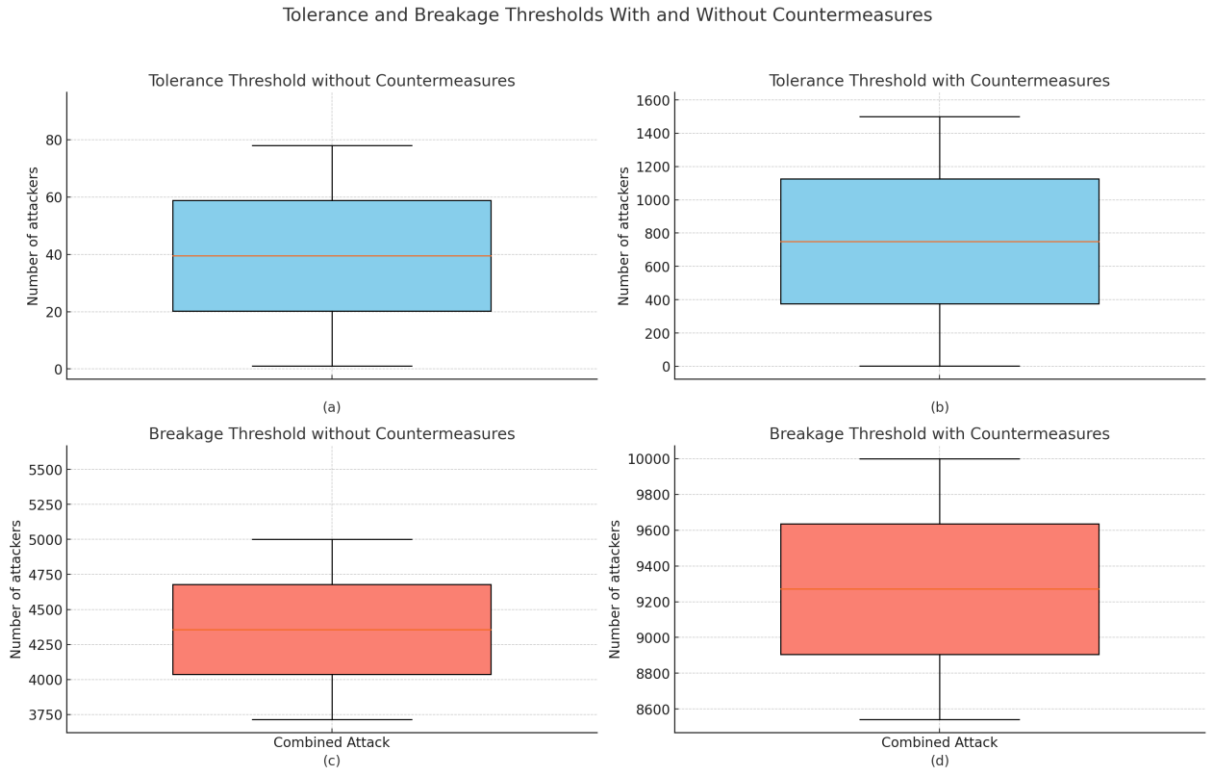

Figure 12. *Tolerance and breakage thresholds in 10,000-node cloud under combined attacks with and without countermeasures.*

Error rates and scalability

Furthermore, we evaluated the performance of the four defence mechanisms by measuring their error rates, indicating detection failures, against various attack scenarios involving cloud sizes ranging from 1000 to 10,000 nodes. Defence 1 is utilized against SYN and UDP flood attacks, defence 2 against HTTP flood, defence 3 against smurf attacks and defence 4 against malformed packet attacks. Error rates were calculated by evaluating detection failures. True positives (*TP*) are correctly identified malicious activities. True negatives (*TN*) are correctly identified legitimate activities. False positives (*FP*) are legitimate activities incorrectly identified as malicious. False negatives (*FN*) are malicious activities incorrectly identified as legitimate. The error rate is computed as $(FP+FN)/(TP+TN+FP+FN)$. This provides a comprehensive measure of detection accuracy.

The experiment results are presented in Figure 13. This figure shows the error rates of the implemented countermeasures for different cloud sizes. The results demonstrate that the error rate of all the defence mechanisms increases proportionally with cloud size. Specifically, for the defence against SYN flood attacks, the error rates were 5.97%, 8.97% and 10.97% in 1000-node, 5000-node and 10,000-node clouds, respectively. Similarly, for UDP flood attacks, the error rates were 5.96%, 8.96% and 10.96%; for HTTP flood attacks, the error rates were 7.7%, 8.7% and 10.7%; for smurf attacks, the error rates were 5%, 6% and 9%; and for malformed packet attacks, the error rates were 4.4%, 8.4% and 13.4%.

Error rates increase with cloud size due to the fact that the number of compromised nodes also increases proportionally. Defence mechanisms often have scalability limits, meaning that their performance degrades as the number of nodes increases. Larger clouds may also experience greater resource constraints, reducing the effectiveness of defence mechanisms as they have fewer resources to analyse and respond to attacks.

The high error rates restrict the detection accuracy, which means that few attacks are being detected; this affects directly the cloud performance making it more vulnerable to attacks, the latter overloads resources and degrades the overall quality of provided service.

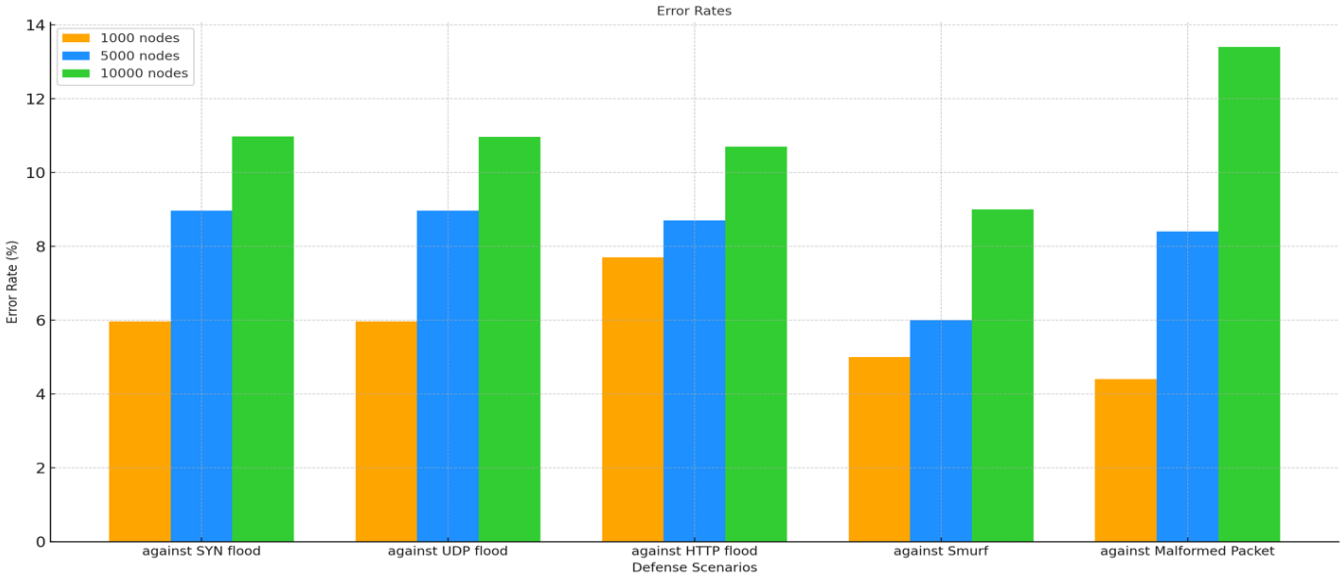

Figure 13. *Error rates of implemented countermeasures in best case in 1000-node, 5000-node and 10,000-node clouds.*

5 CONCLUSION

In this study, we addressed the rising DDoS threats to cloud services by analysing worst-case attack scenarios. Our findings indicate that the cloud maintained its performance even when subjected to multiple simultaneous attacks at tolerance thresholds in all the scenarios. However, our study also highlights that implementing countermeasures may not guarantee improved performance under real-world conditions. Such a study can be useful to those responsible for cloud management. In the absence of countermeasures, or in the case of compromised countermeasures, certain data centres will remain functional under specific conditions. With this in mind, it would also be more appropriate to duplicate data centres carrying out the same tasks.

On the other hand, even in the presence of countermeasures, the cloud can be affected by several simultaneous attacks, leading to its collapse. This is why it would be very useful to provide a second or third line of defence. In our future work, we plan to explore more cloud attacks and enhance existing defences in large-scale environments.

ADDITIONAL INFORMATION AND DECLARATIONS

Conflict of Interests: The authors declare no conflict of interest.

Author Contributions: R.B.: Conceptualization, Writing – Original draft preparation, Supervision, Writing – review and editing, Project administration. F.B.: Methodology, Software, Validation, Writing – Original draft preparation, Visualization, Investigation. A.E.K.: Writing – review and editing.

Statement on the Use of Artificial Intelligence Tools: The authors declare that they didn't use artificial intelligence tools for text or other media generation in this article.

APPENDIX A – GRAPHICAL INTERFACE OF SIMULATION SCENARIOS

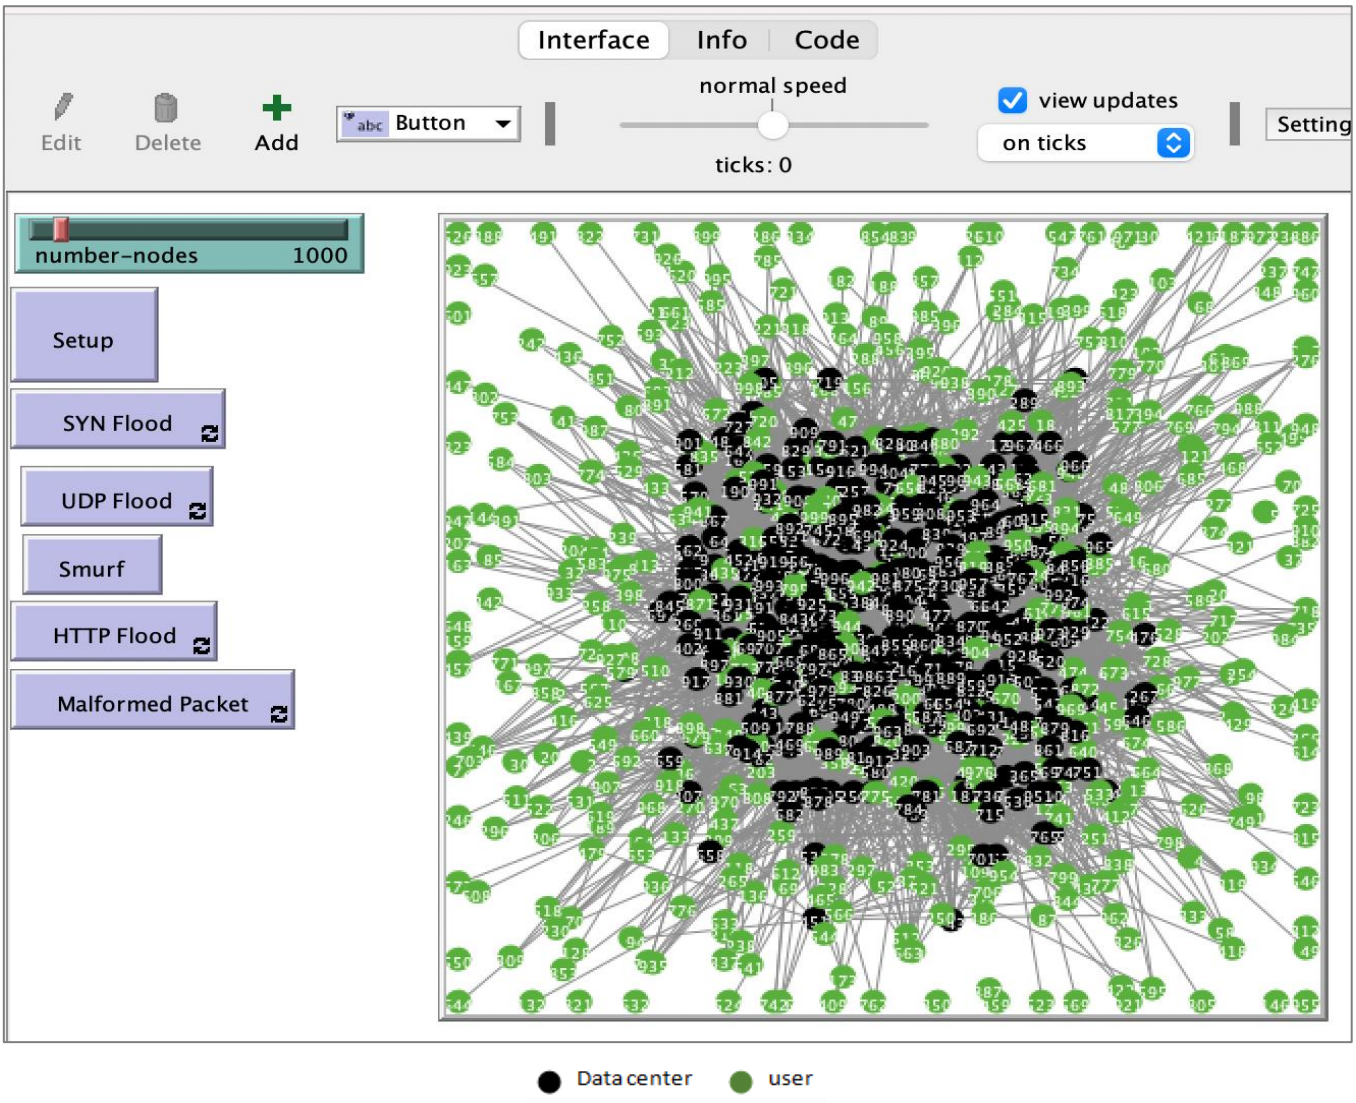

a) Initial state of 1000-node cloud before a SYN flood attack.

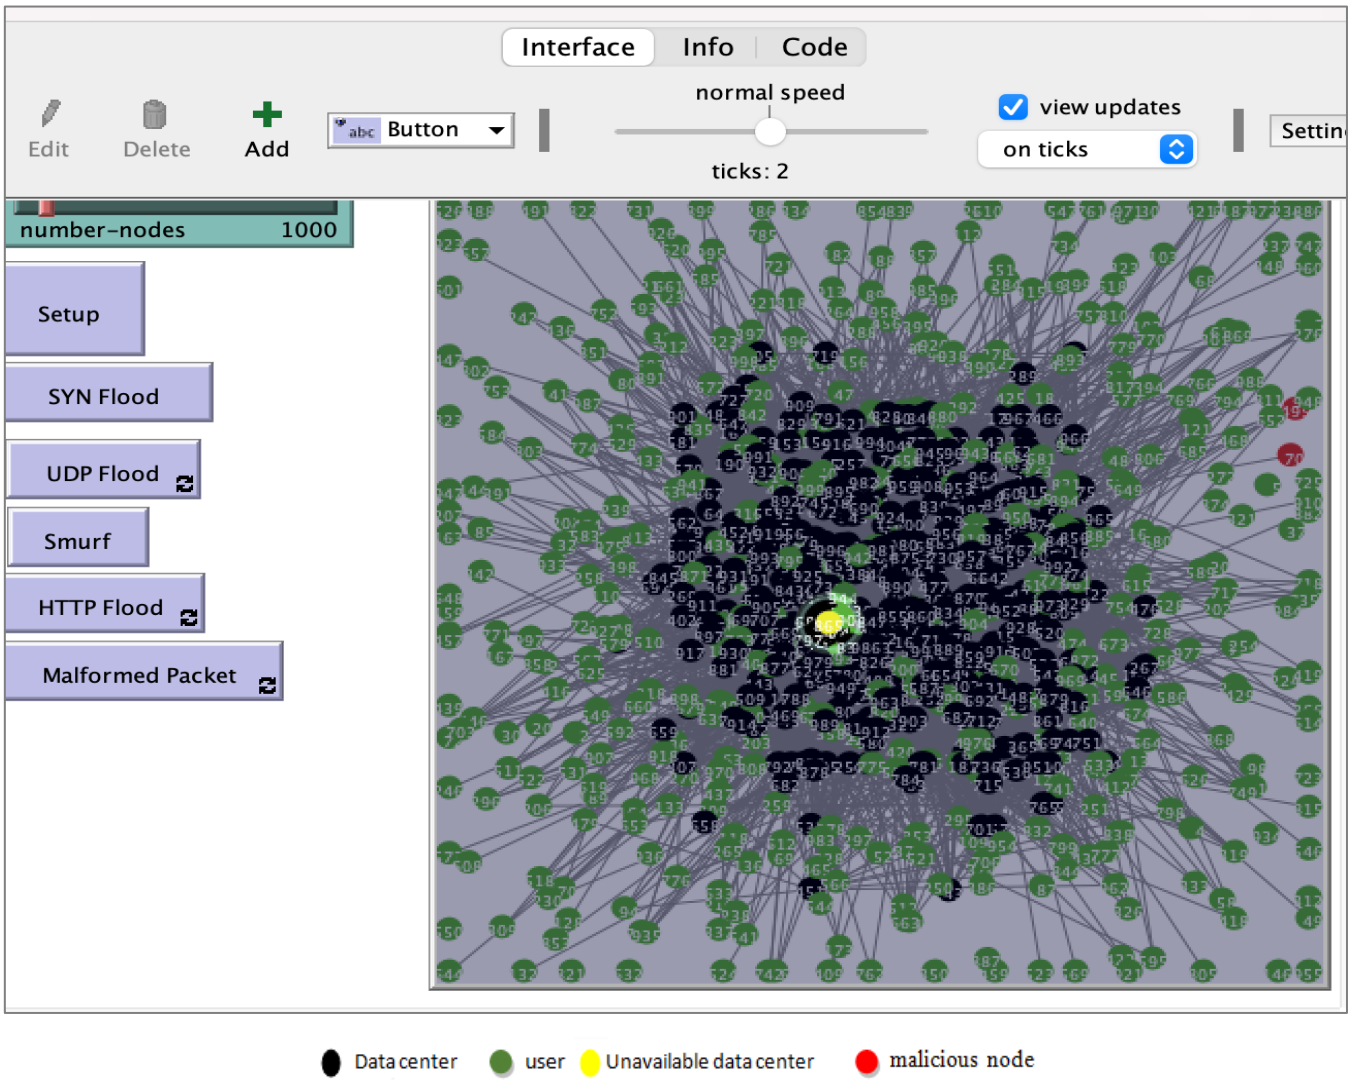

b) After running SYN flood attack (2 malicious nodes,1 unavailable data centre).

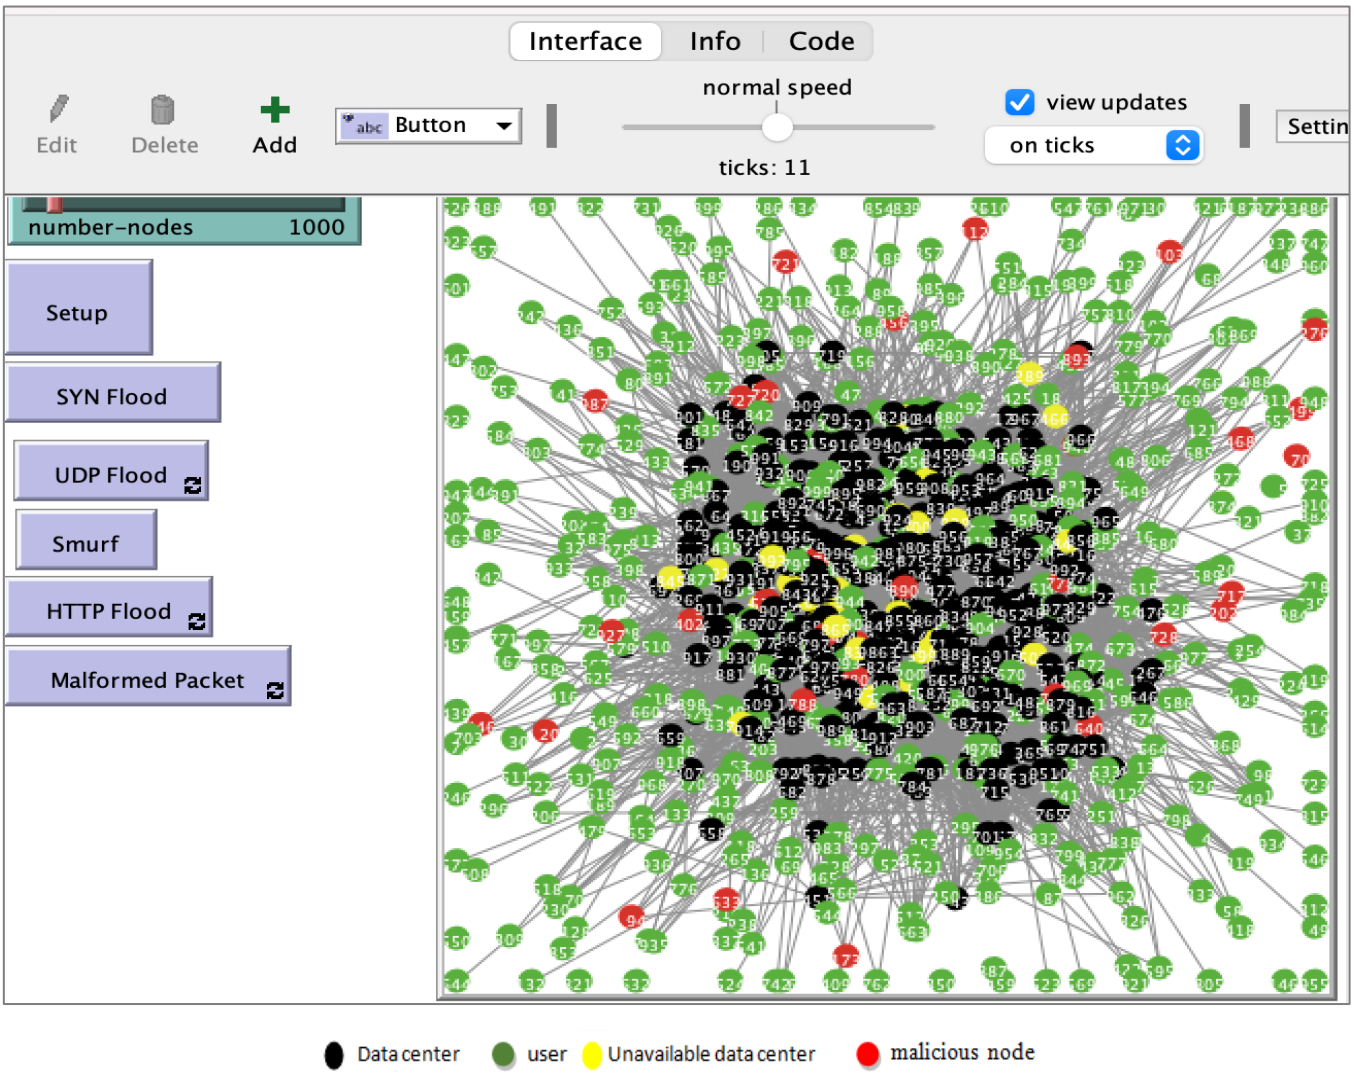

c) After running SYN flood attack (30 malicious nodes, 4% unavailable data centres).

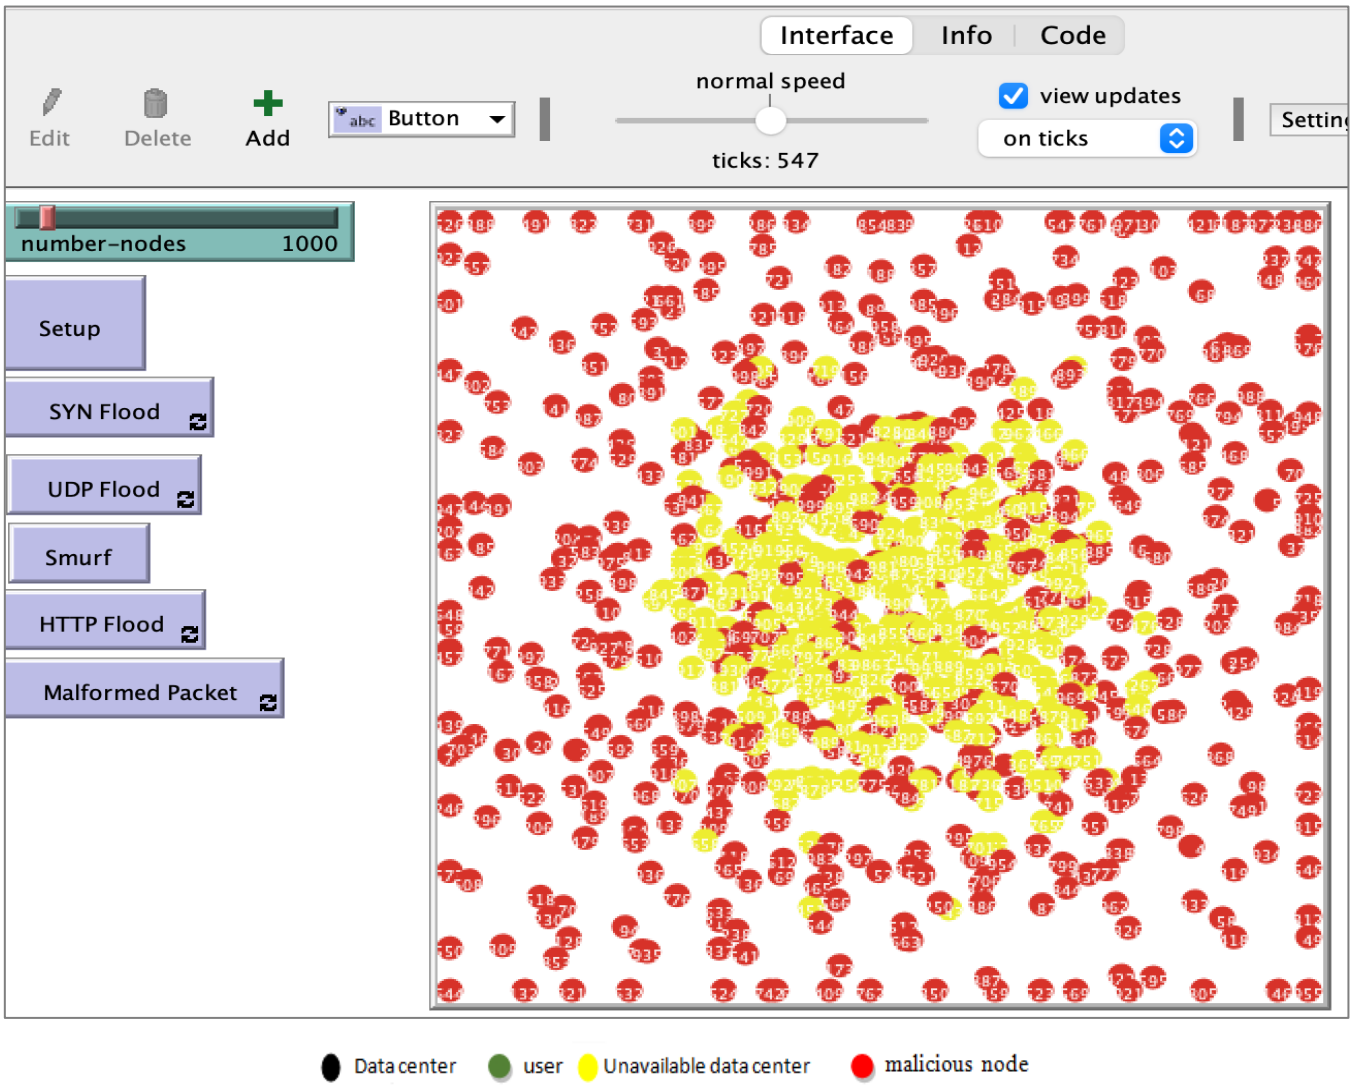

d) Cloud destruction (500 malicious nodes, 100% unavailable data centres).

Figure A1. *Experimental cloud network evolution for SYN flood attacks in 1000-node cloud without countermeasures.*

REFERENCES

Agrawal, N., & Tapaswi, S. (2019). Defense Mechanisms against DDoS Attacks in a Cloud Computing Environment: State-of-the-Art and Research Challenges. *IEEE Communications Surveys and Tutorials*, 21(4), 3769–3795. https://doi.org/10.1109/COMST.2019.2934468

Ahmad, W., Rasool, A., Javed, A. R., Baker, T., & Jalil, Z. (2022). Cyber Security in IoT-Based Cloud Computing: A Comprehensive Survey. *Electronics*, 11(1), 16. https://doi.org/10.3390/electronics11010016

Alosaimi, W., Alshamrani, M., & Al-Begain, K. (2016). Simulation-Based Study of Distributed Denial of Service Attacks Counteract in the Cloud Services. In *2015 9th International Conference on Next Generation Mobile Applications, Services and Technologies*. IEEE. https://doi.org/10.1109/NGMAST.2015.50

Balobaid, A., Alawad, W., & Aljasim, H. (2016). A study on the impacts of DoS and DDoS attacks on cloud and mitigation techniques. In *2016 International Conference on Computing, Analytics and Security Trends (CAST),* (pp. 416–421). IEEE. https://doi.org/10.1109/CAST.2016.7915005

Bouyeddou, B., Harrou, F., Kadri, B., & Sun, Y. (2021). Detecting network cyber-attacks using an integrated statistical approach. *Cluster Computing*, 24(2), 1435–1453. https://doi.org/10.1007/s10586-020-03203-1

Corrêa, J. H., Ciarelli, P. M., Ribeiro, M. R. N., & Villaça, R. S. (2021). Ml-based ddos detection and identification using native cloud telemetry macroscopic monitoring. *Journal of Network and Systems Management*, 29, 1–28. https://doi.org/10.1007/s10922-020-09578-1

Devi, B. S. K., & Subbulakshmi, T. (2021). Cloud DDoS detection and defense system using complex event processing. In *Proceedings - 5th International Conference on Intelligent Computing and Control Systems, ICICCS 2021*, (pp. 118–128). IEEE. https://doi.org/10.1109/ICICCS51141.2021.9432102

Ferretti, L., Magnanini, F., Andreolini, M., & Colajanni, M. (2021). Survivable zero trust for cloud computing environments. *Computers & Security*, 110, 102419. https://doi.org/10.1016/j.cose.2021.102419

Gupta, B. B., & Badve, O. P. (2017). Taxonomy of DoS and DDoS attacks and desirable defense mechanism in a Cloud computing environment. *Neural Computing and Applications*, 28(12), 3655–3682. https://doi.org/10.1007/s00521-016-2317-5

Kanniga Devi, R., Gurusamy, M., & Vijayakumar, P. (2020). An Efficient Cloud Data Center Allocation to the Source of Requests. *Journal of Organizational and End User Computing*, 32(3), 23–36. https://doi.org/10.4018/JOEUC.2020070103

Khandare, H., Jain, S., & Doriya, R. (2023). A Survey on HTTP Flooding—A Distributed Denial of Service Attack. In *Pervasive Computing and Social Networking* (pp. 39–52). Springer. https://doi.org/10.1007/978-981-19-2840-6_4

Liu, Q., & Xing, L. (2021). Survivability and vulnerability analysis of cloud RAID systems under disk faults and attacks. *International Journal of Mathematical, Engineering and Management Sciences*, 6(1), 15. https://doi.org/10.33889/IJMEMS.2021.6.1.003

Mishra, A., Gupta, N., & Gupta, B. B. (2021). Defense mechanisms against DDoS attack based on entropy in SDN-cloud using POX controller. *Telecommunication Systems*, 77(1), 47–62. https://doi.org/10.1007/s11235-020-00747-w

Mthunzi, S. N., & Benkhelifa, E. (2017). Survivability analogy for cloud computing. In *2017 IEEE/ACS 14th International Conference on Computer Systems and Applications (AICCSA),* (pp. 1056–1062). IEEE. https://doi.org/10.1109/AICCSA.2017.219

Nsabimana, T., Bimenyimana, C. I., Odumuyiwa, V., & Hounsou, J. T. (2020). Detection and prevention of criminal attacks in cloud computing using a hybrid intrusion detection systems. In *Intelligent Human Systems Integration 2020, IHSI 2020*, (pp. 667–676). Springer. https://doi.org/10.1007/978-3-030-39512-4_103

Ogwara, N. O., Petrova, K., & Yang, M. L. (2022). Towards the development of a cloud computing intrusion detection framework using an ensemble hybrid feature selection approach. *Journal of Computer Networks and Communications*, 2022, 1–16. https://doi.org/10.1155/2022/5988567

Osanaiye, O., Choo, K. K. R., & Dlodlo, M. (2016). Distributed denial of service (DDoS) resilience in cloud: Review and conceptual cloud DDoS mitigation framework. *Journal of Network and Computer Applications*, 67, 147–165. https://doi.org/10.1016/j.jnca.2016.01.001

Potluri, S., Mangla, M., Satpathy, S., & Mohanty, S. N. (2020). Detection and Prevention Mechanisms for DDoS Attack in Cloud Computing Environment. In *2020 11th International Conference on Computing, Communication and Networking Technologies, ICCCNT 2020*. IEEE. https://doi.org/10.1109/ICCCNT49239.2020.9225396

Raja Sree, T., & Mary Saira Bhanu, S. (2020). Detection of HTTP flooding attacks in cloud using fuzzy bat clustering. *Neural Computing and Applications*, 32(13), 9603–9619. https://doi.org/10.1007/s00521-019-04473-6

Shah, S. Q. A., Khan, F. Z., & Ahmad, M. (2022). Mitigating TCP SYN flooding based EDOS attack in cloud computing environment using binomial distribution in SDN. *Computer Communications*, 182, 198–211. https://doi.org/10.1016/j.comcom.2021.11.008

Sultana, S., Nasrin, S., Lipi, F. K., Hossain, M. A., Sultana, Z., & Jannat, F. (2019). Detecting and Preventing IP Spoofing and Local Area Network Denial (LAND) Attack for Cloud Computing with the Modification of Hop Count Filtering (HCF) Mechanism. In *2019 International Conference on Computer, Communication, Chemical, Materials and Electronic Engineering (IC4ME2).* IEEE. https://doi.org/10.1109/ic4me247184.2019.9036507

Varma, S. A., & Reddy, K. G. (2021). A Review of DDoS Attacks and its Countermeasures in Cloud Computing. In *2021 5th International Conference on Information Systems and Computer Networks, ISCON 2021*. IEEE. https://doi.org/10.1109/ISCON52037.2021.9702388

Velliangiri, S., Karthikeyan, P., & Vinoth Kumar, V. (2021). Detection of distributed denial of service attack in cloud computing using the optimization-based deep networks. *Journal of Experimental and Theoretical Artificial Intelligence*, 33(3), 405–424. https://doi.org/10.1080/0952813X.2020.1744196

Wani, A. R., Rana, Q. P., Saxena, U., & Pandey, N. (2019). Analysis and detection of DDoS attacks on cloud computing environment using machine learning techniques. In *2019 Amity International Conference on Artificial Intelligence (AICAI),* (pp. 870–875). IEEE. https://doi.org/10.1109/AICAI.2019.8701238

Xu, Y., Deng, G., Zhang, T., Qiu, H., & Bao, Y. (2021). Novel denial-of-service attacks against cloud-based multi-robot systems. *Information Sciences*, 576, 329–344. https://doi.org/10.1016/j.ins.2021.06.063

Yan, Q., Yu, F. R., Gong, Q., & Li, J. (2016). Software-defined networking (SDN) and distributed denial of service (DDOS) attacks in cloud computing environments: A survey, some research issues, and challenges. *IEEE Communications Surveys and Tutorials*, 18(1), 602–622. https://doi.org/10.1109/COMST.2015.2487361

Zhao, X. (2017). Study on DDoS attacks based on DPDK in cloud computing. In *2017 3rd International Conference on Computational Intelligence & Communication Technology (CICT)*, (pp. 1–5). IEEE. https://doi.org/10.1109/CIACT.2017.7977325

Acta Informatica Pragensia is published by the Prague University of Economics and Business, Czech Republic | eISSN: 1805-4951