

Minimizing Financial Cost of DDoS Attack Defense in Clouds with Fine-Grained Resource Management

Guttula Sai Teja, Gudelli Venkata, SAI DURGA RAO,

PG scholar, Department of MCA, DNR college, Bhimavaram, Andhra Pradesh.

B.S.MURTHY

(Assistant Professor), Master of Computer Applications, DNR college, Bhimavaram, Andhra Pradesh.

Abstract: As cloud systems grow in popularity, they become more vulnerable to cyber-attacks. A distributed denial of service (DDoS) attack is one of the most notorious cyber-attacks. The attack aims to exhaust the system's resources, rendering it unresponsive to legitimate requests. Around the resource competition, DDoS defence and attack essentially revolve. We have made decisions from a resource management and investment perspective. However, these defense strategies often assume unlimited resources to defend against attacks without considering financial costs. This coarse-grained approach can lead to resource overprovisioning and unnecessary expenses. To tackle this challenge, we perform a thorough analysis and propose a birth-death-based fine-grained resource management system that dynamically scales resources in/out and up/down. This system adaptively chooses the optimal resource leasing mode for cloud service customers, providing cost-effective defense against DDoS attacks. Extensive analyses and experiments based on empirical data validate the efficiency and effectiveness of our approach. Compared to existing methods, our proposal can reduce defence costs by an average of 53.58%, with potential savings of up to 93.75.

Keywords: resource management, cloud security, DDoS attacks

I. INTROUCTION

The aim of the paper is to develop a method for minimizing the financial costs associated with defending against Distributed Denial of Service (DDoS) attacks in cloud environments. This is achieved through the implementation of a fine-grained resource management mechanism that can dynamically and adaptively allocate resources to cloud service customers (CSCs) in response to DDoS threats.

In the past decade, numerous applications have moved to the cloud, driven by its demonstrated advantages. Additionally, the cloud is increasingly embracing new paradigms [1]. More and more

cyber attackers are naturally setting up shop in the cloud. The notorious DDoS attack is a major threat from all the attacks that cause enormous damage to industry [2], [3]. The fundamental issue of DDoS defence and attack revolves around the competition for resources between attackers and defenders, as has been demonstrated. Recently researchers [4], [5], [6] in the battle, the winner is the side that possesses more resources.

In a cloud environment, the two primary entities are cloud service customers (CSCs) and cloud service providers (CSPs). The CSPs control and maintain the cloud infrastructure, whereas DDoS attacks, typically executed through botnets with a few hundred or thousand active bots, frequently fail to impact CSPs significantly. However, individual CSCs, who rely on rented resources from CSPs to run their applications, are more susceptible to these attacks and lack the protection that CSPs inherently provide.

The growth has also made cloud platforms a prime target for DDoS attacks, which can overwhelm resources and disrupt services. Traditional resource management strategies often result in overprovisioning, leading to unnecessary financial burdens for cloud service customers (CSCs).

To address these challenges, this paper introduces a fine-grained resource management mechanism designed to optimize resource allocation during DDoS attacks. By employing a birth-death process model, the proposed approach dynamically adjusts the allocation of virtual machine (VM) instances, ensuring that CSCs can maintain service continuity while minimizing costs. The effectiveness of this method is demonstrated through simulations and empirical data, showing that it can provide a more efficient and cost-effective solution for DDoS defense in cloud environments.

Due to the highly unpredictable and variable of cloud workloads,[9-10] it is essential for cloud service customers (CSCs) to dynamically and adaptively manage resource provisioning to ensure efficient utilization. By the CSP to CSCs the resource provision suggested unfortunately, is either not accurate or blurry. On the basis of performance and experience Most CSCs have to select cloud configurations [11]. To choose the optimal cloud configurations there are many efforts have been made.

Venkataraman *et al.* [12] highlighted the importance of predicting application performance across different resource configurations to find the optimal setup. Alipourfard *et al.* [11] developed CherryPick, a system using Bayesian optimization to identify optimal configurations for recurring big data tasks. These model-based methods rely on pre-trained performance models to guide cloud configuration but are limited by their accuracy, as they are typically trained on non-attack scenario data.

As a result, existing resource provisioning strategies [11], [12], [13] are inadequate for defending against DDoS attacks. Additionally, they often ignore the financial implications for service owners, particularly individual CSCs. Current strategies are typically coarse-grained, like adding virtual machines (VMs) to scale out, which can result in overprovisioning and unnecessary financial costs for CSCs. To address this problem, we introduce a method that lowers the financial impact on CSCs during DDoS defense by dynamically choosing the most efficient resource leasing option.

We introduce a fine-grained resource management mechanism that monitors services in real-time to ensure they perform as expected [15]. Specifically, the mechanism tracks the workloads of the CSC system, adaptively increasing resources during a DDoS attack and reducing them when the attack subsides, ensuring continuous service operation. The numerous existing projects focus on cloud-based DDoS defense [16], [17], [18], most, like the mechanism proposed by Yu *et al.* [8], rely on a coarse-grained approach. Their strategy of simply cloning the same VM(s) to mitigate attacks can lead to resource wastage, as it doesn't

adaptively select different types of VMs for different scenarios.

We argue that the need for resource reinvestment varies depending on the situation, with factors like CPU and memory consumption both playing critical roles during a DDoS attack. Most existing works, including Yu [8], focus only on CPU resource consumption when modelling the system, which can be inefficient if memory is the limiting factor during an attack.

To address this, our approach considers both CPU and memory resources—key factors in most CSC systems and pricing models—when modelling the system. We create a birth-death process model to estimate memory use and processing time based on CPU usage, and then formulates an optimization problem to select the ideal VM type and number for cost-effective attack defense.

The key contributions of the paper are:

- It is crucial to take into account various resource types when safeguarding against DDoS attacks.
- We are creating a model of the birth-death process to gauge memory usage and compute processing duration.
- We are conducting thorough analyses and experiments that demonstrate the effectiveness of our approach and acceptable time overhead.

II. LITERATURE SURVEY

DDoS attacks pose significant threats to internet security by overwhelming a target's services, making them unavailable to legitimate users. These attacks can lead to service disruptions, resource conflicts, and severe financial and reputational damage. While various methods exist to mitigate DDoS attacks, organisations still find it challenging to fully counteract their impact. A highly effective approach combines bandwidth limiting and resource reservations, although these can increase response times for all users during an attack. To address this, we propose a two-tier defence system. The first line separates incoming requests based on the number of connections and

directs them to the second line. Here, bandwidth is limited, and a load balancer directs requests to different containers based on the defence system's decisions. This method efficiently processes legitimate user requests, ensuring service availability for 98% of legitimate traffic, even during intense DDoS attacks. [1]

Edge computing (EC) and container technology enhance computing flexibility for real-time applications but face challenges from low-rate DDoS (LDDoS) attacks that are difficult to detect. Existing defense methods often increase resource consumption and service delays, reducing system efficiency. This paper proposes a solution using Moving Target Defense (MTD) and deep reinforcement learning (DRL) to mitigate LDDoS attacks efficiently. The approach combines lightweight MTD mechanisms and a deep Q-network (DQN) algorithm to balance effectiveness and resource usage. Simulations show significant improvements in security (31.7%) and service quality (26.95%) with minimal overhead, including the lowest response time per request and webpage load time compared to other strategies. [2]

This paper presents a distributed method for detecting DDoS flooding attacks at the traffic-flow level, designed for ISP core networks. The defense system uses distributed change-point detection (DCD) framework with change-aggregation trees (CAT) to identify early traffic anomalies across various network domains. This keeps victim systems from getting too damaged. The system involves attack-transit routers that cooperatively detect and report flooding alerts, which are then aggregated by CAT servers within each ISP domain. These servers collaborate to make final decisions, supported by a Secure Infrastructure Protocol (SIP) to resolve policy conflicts. Simulations on the DETER testbed demonstrate a 98% detection accuracy with only 1% false positives using four network domains, and the system scales effectively to protect up to 84 autonomous system domains, offering broad protection against DDoS attacks.[3]

DDoS attacks are a significant security challenge, especially when executed through distributed botnets. Early detection is crucial to protect users and network infrastructure. This paper introduces Fire Col, a defense system built on intrusion prevention systems (IPSs) deployed at the

ISP level. FireCol creates virtual protection rings around hosts, enabling collaboration between IPSs by sharing traffic information. Extensive simulations and evaluations with real datasets showcase minimal overhead, Fire Col's effectiveness, and its capability for gradual deployment in real networks. [4]

In today's competitive IT landscape, cloud computing provides scalable, on-demand services but faces significant security challenges, particularly DDoS attacks that deplete resources and disrupt service availability. This paper proposes using fog computing, an extension of cloud computing, to analyse and filter DDoS attack traffic at the network edge before it reaches the cloud, enhancing real-time decision-making and reducing data forwarded to the cloud. [5]

The rise in cloud computing has increased vulnerability to DDoS attacks. This paper presents a Multiple Layer Defense (MLD) scheme to detect and mitigate these attacks. MLD operates in two layers: the first sends alarms to cloud management when an attack begins, and the second detects anomalies indicating virtual machines (VM) are compromised. Testing MLD with various DDoS attack ratios shows it effectively reduces energy consumption and SLA violations while maintaining stability. [6]

DDoS attacks in cloud environments are becoming increasingly sophisticated, especially with low-rate DDoS attacks. With the rise of container technology and micro service architecture in cloud computing, which offers more lightweight virtualization and flexible scaling, there's a need to reassess how these new features can defend against DDoS attacks. This paper explores enhancing the resilience of container-based cloud environments against low-rate DDoS attacks by establishing a mathematical model to analyze the feasibility of using these new features. A mitigation strategy is proposed and validated through simulations and experiments. [7]

III. PROPOSED METHOD

Since their invention, cloud services have gained a lot of popularity because they provide

heavy computation and storage spaces at a lower cost. This popularity opened up new ways for cyber attackers to introduce millions of BOTS to spread DDOS attacks, which make cloud servers busy and make genuine requests fail. To address this issue, cloud service providers (CSPs) allocate substantial resources to combat Distributed Denial of Service (DDOS) attacks. However, cloud customers may select resources based on their individual experience and requirements, which may not be optimal and may result in additional costs during DDOS attacks.

To address the aforementioned issue, the author of this paper presents the concept of birth and death, along with fine-grained cloud VM selection. This approach can assist customers in combating DDOS attacks while minimizing costs. In the proposed paper, the author calculates CPU and memory requirements based on request size, increasing BIRTH by 1. Once a request completes, the author decreases BIRTH by 1 and increases death by 1. The author employs a fine-grained algorithm to monitor request requirements. If a request requires heavy resources, the algorithm allocates those resources. If not required, the algorithm releases memory and CPU data to free up resources.

In the proposed paper, the author uses linear programming to identify finer-grained resources, but you ask us to implement a GA algorithm with a dataset. The entire dataset will train the GA algorithm, which will then select an optimal virtual machine (VM) based on the requirements of the request. We will base the GA algorithm on population, mutation, and fitness. We will use the population to generate new random features, then apply mutation to calculate the fitness between the generated and actual data. If the fitness is high, we will select the output; if not, we will continue the process until we achieve an optimised output.

In our proposed work, we are also training the GA algorithm on a cloud VM dataset, and then this algorithm will be applied to test data to predict the optimal VM with less memory and CPU usage. Selecting a VM with less memory and CPU usage will automatically lower the cost.

The proposed method consists of the following key components:

1. **Cloud Users Dataset Analysis:** Analyse datasets of cloud users to understand behaviour, traffic patterns, and resource usage. This involves collecting data, performing behavioural analysis, recognising traffic patterns indicative of DDoS attacks, and examining resource utilisation trends.
2. **DDoS Detection and Classification:** Detect and classify DDoS attacks using anomaly detection algorithms to identify unusual traffic spikes, as well as machine learning classifiers to determine attack type and severity.
3. **Resource Allocation Model:** Create a dynamic resource allocation model based on real-time traffic patterns and attack classifications, including resource pooling, priority assignment, and automatic scaling.
4. **Optimization Using Genetic Algorithms:** Use genetic algorithms to improve the resource allocation model by starting with a population of strategies, testing them with a fitness function, and then using selection, crossover, and mutation to make solutions better over time for cheap DDoS defence.

To train GA algorithm we are using below VM Resources Dataset. In dataset we have information on CPU Usage, Memory Usage, Disk Usage and many more other features. We have evaluated GA VM selection performance in terms of accuracy, precision, recall and FSCORE. In given dataset we have 5 different types of VM's. We have compared GA accuracy with non-GA algorithm accuracy.

We designed modules to implement this project,

- 1) Sign up: user can sign up with the application
- 2) Login: user can login to application
- 3) Load Dataset: after login user can load and process dataset features such as shuffling, normalization etc.

- 4) Minimizing Cloud Cost without GA: in this module user can run V selection algorithm without GA and then calculate selection accuracy
- 5) Minimizing Cost with GA: in this module application will apply GA algorithm for optimal VM selection which can reduce cost
- 6) Predict VM: in this module user can enter REQUEST requirements and then trained GA will predict optimal VM name.

IV. REUSLT

To run the project, install Python 3.7.2, then install the packages listed in the requirements.txt file, and then install MySQL.

Now, double-click on the 'run.bat' file,

```
C:\Windows\system32\cmd.exe
E:\vital\June24\MinimizingCloudCost\python manage.py runserver
C:\Users\Admin\AppData\Local\Programs\Python\Python37\lib\site-packages\pyyaml\_init_.py
C:\Users\Admin\AppData\Local\Programs\Python\Python37\lib\site-packages\pyyaml\_init_.py
Performing system checks...
C:\Users\Admin\AppData\Local\Programs\Python\Python37\lib\site-packages\sklearn\utils\deprecation.py:144: FutureWarning:
The sklearn.metrics.scorer module is deprecated in version 0.22 and will be removed in version 0.24. The corresponding
classes / functions should instead be imported from sklearn.metrics. Anything that cannot be imported from sklearn.metrics
is now part of the private API.
warnings.warn(message, FutureWarning)
C:\Users\Admin\AppData\Local\Programs\Python\Python37\lib\site-packages\sklearn\utils\deprecation.py:144: FutureWarning:
The sklearn.feature_selection.base module is deprecated in version 0.22 and will be removed in version 0.24. The corresponding
classes / functions should instead be imported from sklearn.feature_selection. Anything that cannot be imported
from sklearn.feature_selection is now part of the private API.
warnings.warn(message, FutureWarning)
C:\Users\Admin\AppData\Local\Programs\Python\Python37\lib\site-packages\sklearn\externals\joblib\_init_.py:15: FutureWarning:
sklearn.externals.joblib is deprecated in 0.21 and will be removed in 0.23. Please import this functionality directly
from joblib, which can be installed with: pip install joblib. If this warning is raised when loading pickled models,
you may need to re-serialize those models with scikit-learn 0.21+.
warnings.warn(msg, category=FutureWarning)
System check identified no issues (0 silenced).
You have 15 unapplied migrations(s). Your project may not work properly until you apply the migrations for app(s): admin,
auth, contenttypes, sessions.
Run 'python manage.py migrate' to apply them.
July 06, 2024 - 14:07:36
Django version 2.1.7, using settings 'CloudCost.settings'
Starting development server at http://127.0.0.1:8000/
Quit the server with CTRL-BREAK.
```

Fig. python server

The Python server has started on the screen above. Now, open a browser, enter the URL <http://127.0.0.1:8000/index.html>, and press the enter key to view the page below,



Fig. Sign up Here'

In above screen click on 'Sign up Here' link to get below page,




Fig. user is entering sign up details

In the above screen, the user is entering sign-up details and then pressing the button to complete the sign-up process and get the below page,




Fig. sign up completed

After completing the sign-up process in the above screen, click the 'Login Here' link to access the page below,



Fig. user is login

The user logs in on the screen above, and after logging in, they see the page below,



Fig. user can click on 'Load Dataset' link

In the screen above, the user can click on the 'Load Dataset' link to load and process the dataset, which will result in the page below,



Fig. dataset loaded

In above screen dataset loaded and now click on 'Minimizing Cloud Cost without GA' link to train algorithm without GA to predict optimal VM selection on TEST data and get below prediction accuracy




Fig. without GA algorithm got 97% accuracy on optimal VM selection

On the above screen, without the GA algorithm, I got 97% accuracy on optimal VM selection and could see other metrics like precision, recall, and FSCORE. Click on the 'Minimising Cost with GA'

link to train the algorithm using GA optimisation, resulting in the output shown below,



Fig. GA algorithm got more than 98% accuracy

In above screen with GA algorithm got more than 98% accuracy and can see other metrics also. Now click on 'Predict VM' link to get below page



Fig. 'Request Requirement' data

In above screen paste 'Request Requirement' data which you can copy from 'testData.csv' file from code folder like below screen




Fig. testData.csv file I am selecting one line as request Requirements

In above screen from testData.csv file I am selecting one line as request Requirements and then paste in below screen




Fig. I pasted copied data

In above screen I pasted copied data and then press button to get below page



Fig. blue colour text can see VM0 is the optimal selection

In above screen in blue colour text can see VM0 is the optimal selection for given request and similarly you can copy each line and paste in application to get optimal VM selection



Fig. pasting another requirement

In above screen pasting another requirement and below is the output




Fig. VM3 is selected as Optimal VM

In above screen VM3 is selected as Optimal VM. Similarly you can test other records.

V. CONCLUSION

In this paper, we focus on reducing the financial burden for individual CSCs defending against DDoS attacks. We propose a fine-grained resource management system that dynamically selects the optimal resource leasing mode. Using a birth-and-death process model, we formulate the problem as a constrained optimization challenge and solve it using integer linear programming (ILP). Our approach is validated through detailed analysis and extensive simulations, showing its effectiveness and efficiency with real-world data.

REFERENCES

- [1] Kumar, Anmol, and Mayank Agarwal. "Quick service during DDoS attacks in the container-based cloud environment." *Journal of Network and Computer Applications* 229 (2024): 103946.
- [2] Zhou, Yuyang, Guang Cheng, Zhi Ouyang, and Zongyao Chen. "Resource-Efficient Low-Rate DDoS Mitigation With Moving Target Defense in Edge Clouds." *IEEE Transactions on Network and Service Management* (2024).
- [3] Y. Chen, K. Hwang, and W.-S. Ku, "Collaborative detection of DDoS attacks over multiple network domains," *IEEE Transactions on Parallel and Distributed Systems*, vol. 18, no. 12, pp. 1649–1662, 2007.
- [4] J. Francois, I. Aib, and R. Boutaba, "Firecol: a collaborative protection network for the detection of flooding DDoS attacks," *IEEE/ACM Transactions on Networking*, vol. 20, no. 6, pp. 1828–1841, 2012.
- [5] Bhushan, Kriti. "DDoS attack defense framework for cloud using fog computing." In *2017 2nd IEEE international conference on recent trends in electronics, information & communication technology (RTEICT)*, pp. 534-538. IEEE, 2017.

[6] Khalil, Moataz H., Mohamed Azab, Ashraf Elsayed, Walaa Sheta, Mahmoud Gabr, and Adel S. Elmaghraby. "Maintaining cloud performance under DDOS attacks." *IJCNC* 11, no. 6 (2019): 1-22.

[7] Li, Zhi, Hai Jin, Deqing Zou, and Bin Yuan. "Exploring new opportunities to defeat low-rate DDoS attack in container-based cloud environment." *IEEE Transactions on Parallel and Distributed Systems* 31, no. 3 (2019): 695-706. [

[8] T. Baker, E. Ugljanin, N. Faci, M. Sellami, Z. Maamar, and E. Kajan, "Everything as a resource: Foundations and illustration through internet-of-things," *Computers in Industry*, vol. 94, pp. 62–74, 2018.

[9] "Aliyun suffered from DDoS attack," <http://tech.huanqiu.com/cloud/2014-12/5288347.html>.

[10] "DDoS attack report 2017," <http://cybersecurityventures.com/ddos-attack-report-2017/>.

[11] S. Yu, S. Guo, and I. Stojmenovic, "Can we beat legitimate cyber behavior mimicking attacks from botnets?" in *Proceedings of the 31st IEEE International Conference on Computer Communications*, 2012, pp. 2851–2855.

[12] M. A. Fabian, R. J. Zarfoss, and M. A. Terzis, "My botnet is bigger than yours (maybe, better than yours): why size estimates remain challenging," in *Proceedings of the 1st USENIX Workshop on Hot Topics in Understanding Botnets*, 2007.

[13] S. Yu, Y. Tian, S. Guo, and D. Wu, "Can we beat DDoS attacks in clouds?" *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 9, pp. 2245–2254, 2014.

[14] M. N. Bennani and D. A. Menasce, "Resource allocation for autonomic data centers using analytic performance models," in *Proceedings of the Second International Conference on Autonomic Computing*, 2005, pp. 229–240.

[15] K. Shen, H. Tang, T. Yang, and L. Chu, "Integrated resource management for cluster-based internet services," in *Proceedings of the 5th Symposium on Operating System Design and Implementation*, 2002.