A Hybrid Entropy-Rate Analysis Framework for Lightweight DDoS Detection in Modern Networks

Santhosh K. M.

Lecturer, Department of Computer Engineering, Sree Rama Government Polytechnic College, Thriprayar, Kerala, India
Email: *santhkm[at]gmail.com*

Abstract: ***Distributed Denial-of-Service (DDoS) attacks continue to undermine the stability of modern networks by overwhelming servers and critical services with artificially generated traffic. This work proposes a lightweight hybrid entropy- rate detection framework capable of identifying both high-volume and stealthy low-rate attacks. The approach combines short-term statistical indicators- such as source-IP entropy, packet-rate deviation, byte-volume behavior- and temporal correlation into a unified anomaly-scoring model. By correlating rapid distributional shifts with rate anomalies and temporal persistence, the method enables early detection of diverse attack types while remaining computationally efficient for deployment in resource-constrained institutional networks.***

Keywords: DDoS mitigation, entropy-based analysis, traffic-rate deviation, anomaly detection, lightweight IDS, network security

1. Introduction

The explosive growth of online services, cloud platforms, and interconnected devices has made modern networks highly vulnerable to disruptive cyberattacks. Among these threats, Distributed Denial-of-Service (DDoS) attacks are especially damaging, as they can rapidly exhaust bandwidth, consume server resources, and disrupt essential services. Attackers increasingly employ botnets, IoT device clusters, and adaptive scripts that vary traffic patterns to evade detection, creating significant challenges for conventional defensive mechanisms.

Many traditional detection systems rely on signature matching or static thresholds, which are ineffective against novel attack variants and dynamic traffic patterns. Machine learning approaches improve detection accuracy but often require large labeled datasets and substantial computational resources, limiting their practicality for smaller institutions. There is therefore a critical need for lightweight, reliable detection schemes that can operate in real time on modest hardware.

This paper proposes a hybrid entropy–rate framework that fuses distributional and volumetric indicators into a single anomaly score. Operating on sliding time windows, the system computes source-IP entropy, packet and byte rates, inter-arrival statistics, and temporal correlation measures. These features are combined using weighted indicator functions and adaptive thresholds to detect both aggressive floods and subtle low-rate attacks without heavy processing overhead. Prior studies also support entropy-based anomaly detection [1]. Rate-based deviation methods have been widely used for DDoS detection [2].

2. Literature Review

Early DDoS detection work emphasized signature-based IDS, which perform well against known attacks but cannot generalize to novel threats. As a result, statistical and behavior-based approaches gained traction, offering the ability to detect anomalies without prior knowledge of attack signatures.

Entropy-based metrics measure the randomness of traffic distributions and are effective at highlighting concentration of traffic among a small set of hosts- an indicator of coordinated attacks. However, entropy can be insensitive to carefully crafted low-rate attacks that preserve distributional uniformity. Rate-based techniques, on the other hand, detect volumetric anomalies but may miss stealthier behaviors.

Machine learning and deep learning methods have demonstrated strong classification performance in controlled environments, but their requirements for training data, tuning, and computational resources make them less suitable for real-time deployment on constrained devices. Recent research therefore explores hybrid approaches that combine entropy and rate features, sometimes augmented by temporal analysis, to achieve robust detection without heavy models. Prior studies also support entropy-based anomaly detection [1]. Similar hybrid entropy–rate strategies are discussed in earlier research [3], [4].

Nevertheless, many hybrids still rely on static thresholds or complex fusion strategies that complicate deployment. This paper addresses these gaps by proposing an adaptive, low-overhead hybrid scoring mechanism that balances sensitivity and robustness for institutional environments.

3. System Architecture

The proposed framework is organized as a modular pipeline optimized for efficiency and clarity. It comprises four principal components: Traffic Capture, Feature Extraction, Hybrid Scoring Engine, and Decision Layer. Each module is deliberately lightweight to facilitate operation on gateway devices and edge routers.

a) **Traffic Capture Layer:** This module aggregates packets into short sliding windows (typically one second), extracting essential metadata such as source IP, destination IP, packet size, timestamps, and protocol flags. The sliding-window approach balances temporal resolution and processing overhead.
b) **Feature Extraction Module:** For each window, compact statistical features are computed, including source-IP Shannon entropy, packet and byte rates, inter-arrival time

variance, and simple port distribution statistics. These features capture both distributional and volumetric aspects of traffic. Prior studies also support entropy-based anomaly detection [1]. Rate-based deviation methods have been widely used for DDoS detection [2].

c) **Hybrid Scoring Engine:** The extracted metrics are compared against dynamically maintained baselines. Deviations are measured relative to mean and standard deviation (e.g., $\mu \pm k\sigma$) and combined via weighted indicator functions to form an overall anomaly score S.

d) **Decision Layer:** When the anomaly score exceeds a configured alert threshold, the system emits an alert, logs diagnostic metadata, and can trigger automated mitigation hooks. The decision logic is intentionally transparent and deterministic for easy operational interpretation.

4. Implementation Methodology

The implementation adopts a sliding-window segmentation with windows of fixed duration to capture micro-level traffic dynamics while keeping processing overhead low. Aggregation windows simplify statistical calculations and reduce memory requirements.

Within each window, the system computes several lightweight features: Shannon entropy over source-IP frequency, packet count and byte totals, inter-arrival time variance, and a temporal correlation index comparing current and previous windows.

Baselines for each metric are constructed during benign operation, and dynamic thresholding is applied using statistical deviation rules such as entropy < (mean − 2σ) or rate > (mean + 3σ). These adaptive checks reduce false positives during legitimate traffic surges.

The hybrid anomaly score S is computed by summing weighted indicator outputs for each feature, with tunable weights α, β, γ. The final classification is a simple threshold comparison against T_alert, ensuring fast, explainable decisions without requiring model training.

Overall, the methodology prioritizes computational simplicity and deterministic behavior, enabling the detector to run on commodity hardware with minimal resource consumption.

4.1 Feature Computation Equations

The following lightweight statistical indicators are computed within each sliding window:
Entropy (distributional variation): $H(W) = - \Sigma\, p(i) \log_2 p(i)$
Packet rate: $R(W) = N / T$
Byte rate: $B(W) = \Sigma(\text{size}) / T$
Temporal correlation: compares current-window features with historical windows to detect persistent abnormality.

5. Detection Algorithm

In order to operationalize the hybrid detection framework, the extracted features are processed by a deterministic routine that produces a binary decision at the end of every window. T_alert denotes the global anomaly threshold; in this work, T_alert = 2 is used to balance sensitivity and false positives.

```
Algorithm 1: Hybrid Entropy–Rate Detection
Input: Live traffic stream T (sliding windows of 1s)
Output: DDoS Alert / Normal
1) Capture all packets in current window W
2) Compute source-IP distribution p(i) and Shannon
   entropy $H(W) = -\Sigma\, p(i) \log_2 p(i)$
3) Compute packet rate $R(W) = N / \text{window\_duration}$ and
   byte rate $B(W)$
4) Compute inter-arrival variance $\sigma_{ia}(W)$ and temporal
   correlation CT
5) Compute hybrid score $S = \alpha \cdot I[R>\mu R+3\sigma R] + \beta \cdot I[H<\mu H-2\sigma H] + \gamma \cdot I[CT>\tau]$
6) If $S \geq$ T_alert then: raise DDoS alert, log flow metadata,
   and trigger mitigation
7) Else: label window as NORMAL and continue
   monitoring
```

Notes: I[cond] is indicator function; α,β,γ are weights (default=1); μ and σ derived from baseline.

6. Results and Evaluation

To validate the proposed detector, we conducted controlled experiments varying attacker counts and background traffic. Performance is evaluated using detection accuracy, false positive rate, precision and recall, and computational scalability. Attackers were simulated with counts ranging from 10 to 50 while background traffic mimicked typical institutional usage patterns.

6.1 Detection Accuracy

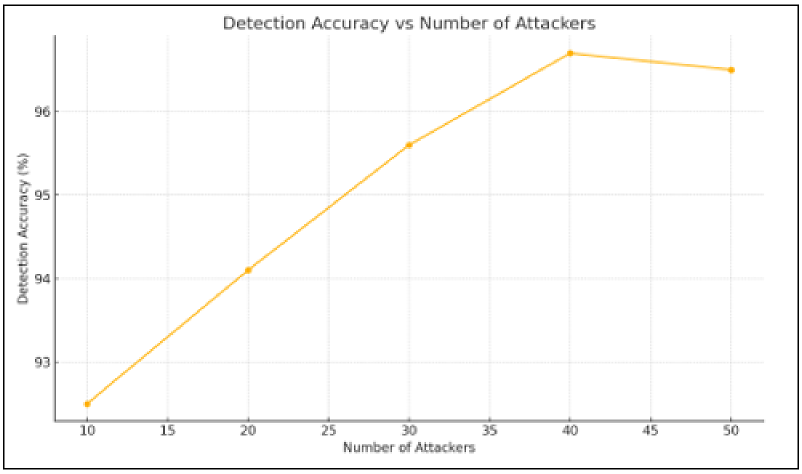

Figure 1: Detection accuracy under varying attacker loads

Detection accuracy remains high across attacker intensities, reflecting the hybrid model's ability to combine complementary cues. Fig. 1 shows that accuracy improves as attacker populations grow because volumetric and distributional signals reinforce each other.

6.2 False Positive Rate

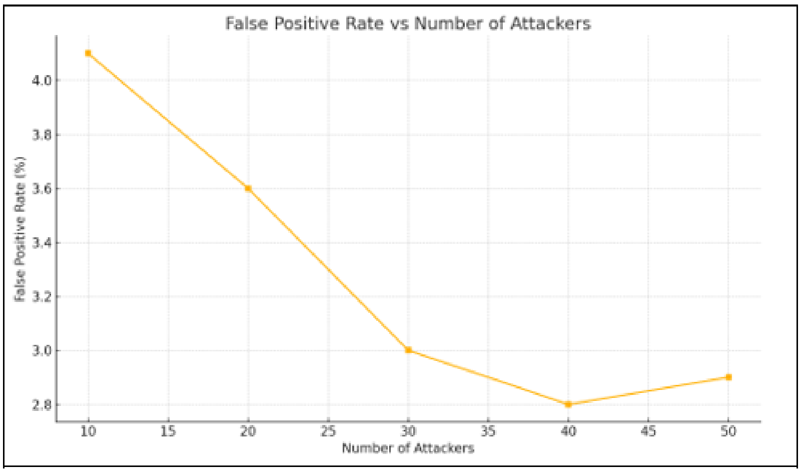

Figure 2: False positive rate across simulations.

False positive rate declines as adaptive thresholds stabilize around normal traffic behavior; temporal correlation prevents transient bursts from causing alerts, as illustrated in Fig. 2.

6.3 Precision and Recall

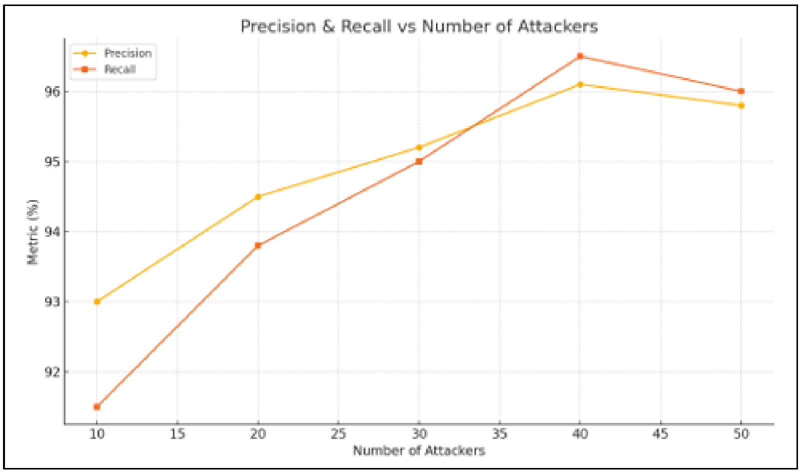

Figure 3: Precision and recall maintain a favorable balance, indicating robust detection.

Precision and recall maintain a favorable balance (Fig. 3), meaning the detector identifies most attack windows while producing few false alarms.

6.4 Scalability

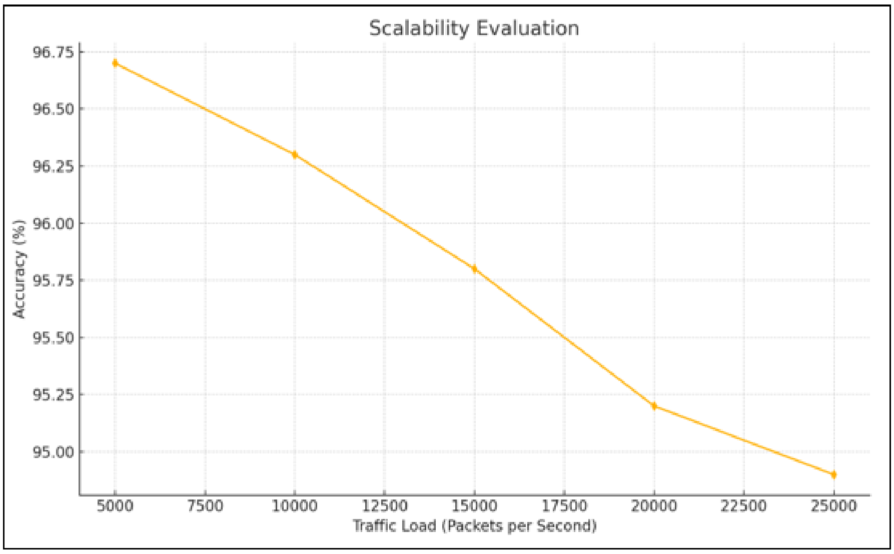

Figure 4: Scalability analysis with increasing packet load.

Scalability tests reveal that the detector processes windows with linear complexity and maintains performance up to high packet rates (Fig. 4), indicating suitability for deployment on modest hardware.

7. Conclusion

This study presented a practical and lightweight hybrid detection framework that integrates entropy variations, rate deviations, and temporal characteristics to identify DDoS attacks. The approach emphasizes low computational overhead and high interpretability, making it well-suited for institutional deployments. Experiments demonstrate robust accuracy, low false positives, and scalable performance. Future work includes integrating the detector with SDN controllers for automated mitigation, extending baselines with long-term learning, and testing in production network environments.

References

[1] Y. Zhang et al., “Entropy-based anomaly detection in high-speed networks,” IEEE Access, 2021.
[2] M. Alzahrani, “Hybrid intrusion detection systems,” Int. J. Comput. Netw. Commun., 2020.
[3] S. Rao et al., “Deep learning approaches for DDoS detection,” Computer Communications, 2022.
[4] J. Li and X. Wu, “Traffic entropy metrics for flood detection,” IEEE Trans. Netw. Serv. Manag., 2019.
[5] J. Kim et al., “Flow-based lightweight IDS framework,” Sensors, 2021.