Transaction Analysis of Russian Malicious Cyber Activity on SINET

Ruo Ando[1†],

2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430

Abstract
We address the problem of detecting malicious communications in environments where IP-based reputation is costly to maintain and fragile under adversarial churn. We propose a transaction-centric representation that groups related log events into sessions using standard maxspan and maxpause constraints, and we extract three interpretable features per transaction—event count, duration, and concurrency. Embedding transactions in this three-dimensional (3D) space yields geometry that is both human-interpretable and amenable to lightweight outlier detection; in practice, suspicious behaviors manifest as stable high-density departures across inbound/outbound traffic and parameter regimes. The approach is payload-agnostic (robust to encryption), reduces dependence on external threat intelligence, and lowers downstream learning complexity by working in a low-dimensional, well-separated feature space. We deploy the method on SINET, a large Japanese academic backbone with dynamic addressing and heavy international connectivity, and we demonstrate that (i) transaction geometry reveals characteristic differences between maxspan- and maxpause-driven sessionization, (ii) outlier regions identified in the 3D space align with operator-validated anomalies across directions of flow, and (iii) simple density/thresholding schemes operating on these features provide an effective screening layer that complements conventional reputation pipelines. Collectively, our results indicate that transaction-based modeling offers a practical, computationally economical alternative for first-line malicious-communication detection in high-throughput research networks.

Keywords:
Transaction Analysi, Russian Malicious Cyber Activit, SINET, High through

1. Introduction

For cybersecurity reputation analysis to be effective, reliable data sources are essential. In particular, leveraging threat intelligence to collect real-time information on malware, phishing sites, and malicious IP addresses is crucial. In addition, diverse data such as network traffic, authentication logs, DNS queries, and firewall logs must be gathered and analyzed. Integrating external sources like VirusTotal, Shodan, AbuseIPDB, and AlienVault OTX further enhances coverage.

Accuracy and freshness are critical because reputation information changes over time, requiring real-time or frequent updates. False positives and false negatives must be minimized; appropriate filtering and whitelisting improve reliability. Past misclassifications should be used to retrain models and improve precision.

Analyzing Russia's cyber activities is challenging for several reasons. First, state involvement and limited transparency hinder reputation analysis. Many operations appear state-sponsored, with agencies such as the FSB and GRU playing active roles. As a result, attackers and infrastructure may be protected by the government, complicating data collection by external threat-intelligence and research organizations. Moreover, Russia's legal framework restricts extradition and limits international cooperation in cybercrime investigations.

Second, Russian cyber actors employ advanced anonymization techniques that complicate tracking. Tor, VPNs, proxy chains, and Fast Flux DNS are frequently used to obfuscate origins. Fast Flux DNS, in particular, rotates domain associations rapidly, degrading the effectiveness of domain-reputation analysis. By leveraging cloud services and compromised servers, attackers can further disguise their locations, frustrating reputation-based tracking.

Manuscript received September 5, 2025
Manuscript revised September 20, 2025
https://doi.org/10.22937/IJCSNS.2025.25.9.1

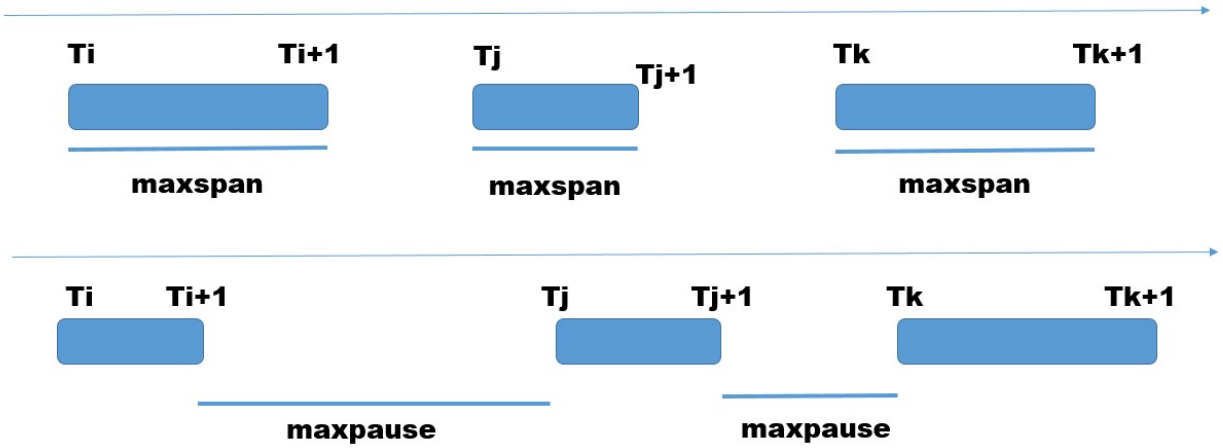

Figure 1. Transaction analysis

Third, groups such as APT28 and APT29 commonly employ disposable infrastructure and zero-day exploits. IP addresses and domains are often rotated per campaign, rendering traditional blacklists ineffective. Zero-day vulnerabilities allow bypassing of existing security databases, making detection more difficult. This dynamic, rapidly evolving strategy undermines conventional reputation techniques and necessitates more adaptive, real-time intelligence. We have addressed two challenges on SINET, a large-scale academic backbone network, as follows:

Large IP address space and anonymity on SINET: SINET utilizes an extensive IP space, much of it dynamically assigned, which complicates reputation-based evaluation compared with corporate networks that have stable allocations. For privacy protection, some communications are anonymized, making it harder to trace origins and to distinguish legitimate research from potential abuse.

International connectivity and external data exchange: SINET connects research institutions worldwide and facilitates large-scale data exchange. Academic networks prioritize openness, leading to frequent international interactions. Universities and research centers often use international address space and cloud services, which can reduce the fidelity of external reputation databases. Sudden traffic surges from specific countries—suspicious in corporate settings—are commonplace in academia, increasing the risk of misinterpretation by traditional methods.

2. Reputation analysis

The reputation score. R_j of an IP address IP_j is computed from attributes X_ij as follows:

$$R_j = \sum_{i=1}^{n} w_i X_{ij}, \qquad (1)$$

where:

- j indexes the target IP address
- i indexes attributes (e.g., attack frequency, blacklist count)
- X_ij is the value of attribute i for IP j
- w_i is the weight of attribute i
- n is the number of attributes

Typical attributes. X_ij includes:

- Blacklist entries: number of times the IP is flagged as malicious;
- Malware detections: incidents attributed to the IP;
- Anomalous traffic: high-volume or suspicious patterns (e.g., DDoS);
- Geographical risk: risk associated with the IP's location;
- Open ports: exposure of unnecessary services;

Protocol anomalies: unauthorized protocol/port usage.

Normalization. Attributes are normalized to [0,1]:

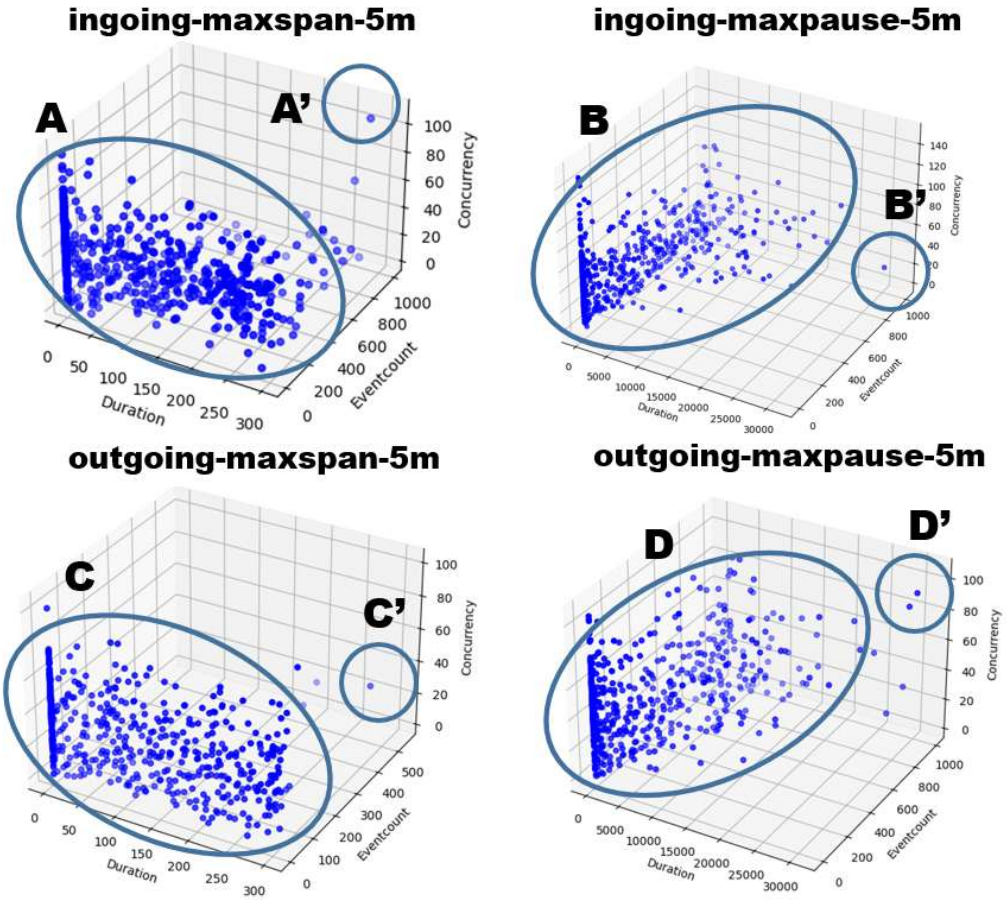

Figure 2. 3D plot of duration, concurrency and event count

$$X'_{ij} = \frac{X_{ij} - \min(X_i)}{\max(X_i) - \min(X_i)}, \qquad (2)$$

where X'_{ij} is the normalized value.

Risk classification. Given thresholds T1 and T2,

$$\text{Risk} = \begin{cases} \text{High}, & R_j > T_2, \\ \text{Medium}, & T_1 \le R_j \le T_2, \\ \text{Low}, & R_j < T_1. \end{cases} \qquad (3)$$

Temporal updates. Because an IP's reputation changes over time, we use exponential smoothing:

$$R_{j,t} = \alpha R_{j,t-1} + (1-\alpha) \cdot \text{NewScore}, \qquad (4)$$

where 0 <= α <= 1.

3. Transaction analysis

Splunk's transaction command groups multiple log events into a single transaction representing a sequence of operations—for example, interactions across systems or a user session. It supports:

- Grouping related events by a shared field (e.g., IP address, user ID);
- maxspan: the maximum total duration of a transaction;
- maxpause: the maximum allowed idle time between consecutive events;
- Retrieval of detailed per-transaction event lists and timestamps;
- Pattern/anomaly discovery for security and operations.

Figure1 illustrates two transaction settings. The upper panel shows a maxspan-based analysis, which

constrains the total span considered a single transaction (e.g., per IP). Larger maxspan yields larger within-transaction metrics such as event count and concurrency, while smaller maxspan increases the number of transactions.

The lower panel shows a maxpause-based analysis, which constrains the permitted idle time between events: if the gap exceeds maxpause (e.g., 5 minutes), a new transaction begins.

4. Experiment

Duration. Duration is the elapsed time from a transaction's first to last event:

$$\text{Duration} = t_{\text{end}} - t_{\text{start}}.$$

In Splunk, this is the field \texttt{duration}.

Event count. Event count is the number of events in a transaction:

$$\text{EventCount} = N,$$

is represented as event_count.

Concurrency. Concurrency measures how many events are active simultaneously within a transaction and is useful when transactions overlap:

$$\text{Concurrency} = \frac{\sum_{i=1}^{N} \text{ActiveEvents}(t_i)}{\text{Duration}}.$$

Results. These metrics are readily computed using transaction. For example:

```
index=your_index sourcetype=your_sourcetype
| transaction startswith="start_condition"
  endswith="end_condition"
| table eventcount concurrency duration
```

Here, startswith and endswith define the conditions that mark the beginning and end of a transaction.

Figure 3 analyzes Russian traffic on SINET, split into inbound and outbound, and into two transaction modes—maxspan-based and maxpause-based—yielding four 3D plots. The maxspan-based mode generally exhibits lower concurrency, whereas the maxpause-based mode tends to yield lower durations. In maxspan-based analysis, unusually high concurrency indicates anomalies.

In the top-left (inbound, maxspan-based), Area A shows high concurrency, duration, and event count. In outbound maxspan-based analysis, Area C shows high duration and event count; a concurrency of 23 is anomalous. In the maxpause-based analysis (bottom-right), Area D stands out with high values for all three metrics. The top-right (Area B) has high duration and event count but low concurrency; these are still flagged as anomalies when thresholds are exceeded.

With maxspan set to 5 minutes, points cluster at low event counts with wide spread along the duration axis. With maxpause at 5 minutes, clusters concentrate at low durations with broader spread along the event-count axis. Outliers for inbound maxspan and outbound maxpause appear in similar regions—where all three metrics are high—while outliers in Areas B and C occur at high duration and event count but low concurrency.

5. Discussion

Transaction geometry and empirical findings.
Modeling security-relevant behavior as transactions yields a three-dimensional (3D) geometry—event count, duration, concurrency—that is both interpretable and discriminative under encryption. On SINET, inbound and outbound traffic exhibit more similar distributions in this space than in conventional flow counters, suggesting that sessionization imposes a unifying structure on heterogeneous workloads. Maxspan-based sessionization expands duration and occasionally raises concurrency, while maxpause-based sessionization fragments idle-heavy behaviors and concentrates near short durations; outliers differ accordingly (high concurrency under maxspan, bursty high event counts under maxpause).

Sessionization sensitivity.
Session boundaries are the principal inductive bias. Overly large maxspan can merge unrelated workflows; overly small maxpause can fragment coherent ones. Practical tuning combines per-sourcetype heuristics (e.g., keepalive intervals) with small grid scans that maximize separation in the 3D features. A multi-scale strategy—evaluating short/medium/long parameter pairs in parallel—stabilizes outlier consensus and mitigates single-scale bias.

Integration with reputation and evaluation.
Reputation and transaction geometry are complementary: the latter provides payload-agnostic, feed-light screening; the former adds precision when high-quality labels exist. A staged pipeline first applies geometric screening to preserve recall and reduce candidate volume, then routes survivors to reputation checks and content-aware analytics (e.g., passive DNS/TLS fingerprints). Recommended reporting includes screening reduction at fixed recall, mean time-to-flag, and stability of outlier sets across sessionization regimes.

Scalability and operations on SINET.
Computing the three features requires only per-transaction aggregation with state proportional to open transactions per key. Sliding-window operation, coarse timestamp bucketing, and pre-partitioning by direction or high-variance keys bound memory and compute. Emitting compact records eventcount, duration, concurrency) enables lightweight density or nearest-neighbor scoring and supports analyst workflows via stratified sampling of outlier strata.

Privacy and adversarial considerations.
Because the method is payload-agnostic, it reduces exposure to sensitive content; nonetheless, analyst-facing views should hash or aggregate identifiers, enforce role-based access, and adopt retention limits. Adaptive adversaries may throttle to remain in low-concurrency/low-duration regions or inject chaff; multi-scale sessionization and cross-feature constraints raise evasion cost, while fusing side-channel signals (e.g., JA3/JA4, DNS volatility) further hardens screening. Finally, we analyzed the questionnaires data and based on the findings, some results and recommendations are suggested. The next section shows the results of the companies' survey.

6. Related work

Research on reputation analysis includes Jøsang et al. [1], which surveys reputation-based systems. Chandola et al. [2] discuss anomaly-detection techniques in the financial domain. For transaction-based anomaly detection, Han and Kamber [3] describe clustering methods in distributed systems, and Hosseini and Buyya [4] analyze malicious transactions in enterprise networks.

Regarding SINET, Aoyama et al. [5] evaluate high speed data transfer in SINET5, and Kitagawa et al. [6] study security enhancements for high-speed academic networks. Yu and Zhang [7] review deep-learning approaches to network anomaly detection, while Krishnamurthy et al. [8] present an autoencoder-based method that improves detection accuracy for transaction analysis. For visualization, Lakkaraju and Yurcik [9] and Münz and Li [10] explore traffic visualization techniques, including 3D plots for forensic analysis.

From an operational standpoint, a hybrid of transaction analysis and machine-learning analysis is the most pragmatic choice.

7. Conclusions

We introduced a transaction-centric approach for detecting malicious communications that operates on three interpretable features—event count, duration, and concurrency—derived from sessions constructed via maxspan and maxpause. By embedding transactions in a three-dimensional (3D) space, our method provides human-interpretable geometry that enables lightweight screening with simple density- or threshold-based criteria, while remaining payload-agnostic and reducing dependence on external reputation feeds.

Deployed on SINET, a large academic backbone characterized by dynamic addressing and substantial international connectivity, the approach revealed consistent structure across inbound/outbound directions and sessionization regimes. In particular, (i) maxspan-based sessionization tends to yield lower concurrency and wider dispersion along duration, whereas (ii) maxpause-based sessionization concentrates near short durations with broader spread along event count. Outlier regions identified in this 3D space aligned with operator-validated anomalies, including cases exhibiting simultaneously high concurrency, duration, and event count. These observations indicate that the transaction geometry is stable enough to serve as an effective first-line filter that complements conventional IP-reputation pipelines and prioritizes analyst attention.

Limitations.

Our results depend on the quality of session boundaries: inappropriate maxspan/maxpause settings may fragment or merge behaviors and distort feature geometry. Workload composition (e.g., scheduled data transfers, teaching periods), diurnal effects, and bursty research traffic can also shift distributions. Ground-truth labels at backbone scale are inherently sparse, complicating quantitative benchmarking. Finally, adaptive adversaries might attempt to inject chaff or manipulate inter-event timing to evade detection.

Implications.

The low-dimensional, interpretable representation offers a pragmatic path to deployable screening at scale, particularly where encrypted payloads limit deep inspection. Because the features are simple and separable, the approach can reduce the computational burden of downstream learning and shorten triage cycles for security operations in research networks.

We plan to (i) automate parameter selection for sessionization via adaptive/online estimation, (ii) incorporate multi-scale transaction modeling and temporal drift handling, (iii) fuse transaction features with passive DNS/TLS/flow metadata and existing reputation signals, (iv) conduct comparative studies against LOF/DBSCAN/Isolation Forest under controlled replay and adversarial scenarios, and (v) prototype a streaming implementation for line-rate operation with operator-in-the-loop feedback. Broader validation on additional academic and enterprise backbones will further assess generalizability and robustness.

References

[1] Jøsang, A., Ismail, R., and Boyd, C. (2007). “Reputation-based systems: A survey and taxonomy.” Computers & Security.

[2] Chandola, V., Banerjee, A., and Kumar, V. (2009). “Anomaly detection: A survey.” ACM Computing Surveys.

[3] Han, J., and Kamber, M. (2016). “Anomaly detection using transaction-data clustering in distributed systems.” Elsevier.

[4] Hosseini, S. A., and Buyya, R. (2018). “Dynamic analysis of malicious transactions in enterprise networks.” IEEE Transactions on Network and Service Management.

[5] Aoyama, T., et al. (2020). “High-speed data transfer in academic networks: Case studies on SINET5.” Journal of Information Processing Society of Japan.

[6] Kitagawa, H., et al. (2019). “Advanced security techniques in high-speed academic networks.” IEICE Transactions on Communications.

[7] Yu, H., and Zhang, D. (2020). “Deep learning-based network anomaly detection: A survey.” ACM Computing Surveys.

[8] Krishnamurthy, P., et al. (2021). “Autoencoder-based transaction anomaly detection for malicious communications.” Proceedings of IEEE Big Data.

[9] Lakkaraju, K., and Yurcik, W. (2015). “Visualization techniques for network security and anomaly detection.” ACM SIGCOMM Workshop.

[10] Münz, A., and Li, S. (2017). “3D visualization of network traffic for forensic analysis.” IEEE Visualization Symposium.

Ruo Ando received Ph.D. from Keio University in2006. He is now associate professor by specialappointment of National Institute of Informatics
Since 2016. Before joining NII, he worked as seniorresearcher of National Institute of Information andCommunications Technology since 2006. His researchinterests focus on network security, informationsecurity and big data mining technologies. He received OutstandingLeadership Award in the 8th IEEE International Conference on Dependable,Autonomic and Secure Computing (DASC-09) at China in 2009. He is themember of Trusted Computing Group JRF (Japan Regional Forum) in 2008-2015. He worked in project "Next Generation Security Info-Security R&D"METI (FY2008-10). He was engaged in project "Unknown malwaredetection using incremental malware detection" MEXT FY(2012-2015).